arXiv:1405.5631v1 [math.RA] 22 May 2014

Towards an efficient MEAT-AXE algorithm using f-cyclic matrices: the density of uncyclic matrices in $\mathrm{M}(n, q)$

To John Cannon and Derek Holt in recognition
of their distinguished contributions to mathematics,
and in particular, to computation and the Magma system.

S. P. GLASBY AND CHERYL E. PRAEGER

ABSTRACT. An element X in the algebra $\mathrm{M}(n, \mathbb{F})$ of all $n \times n$ matrices over a field $\mathbb{F}$ is said to be f-cyclic if the underlying vector space considered as an $\mathbb{F}[X]$-module has at least one cyclic primary component. These are the matrices considered to be "good" in the Holt-Rees version of Norton's irreducibility test in the MEAT-AXE algorithm. We prove that, for any finite field $\mathbb{F}_q$, the proportion of matrices in $\mathrm{M}(n, \mathbb{F}_q)$ that are "not good" decays exponentially to zero as the dimension n approaches infinity. Turning this around, we prove that the density of "good" matrices in $\mathrm{M}(n, \mathbb{F}_q)$ for the MEAT-AXE depends on the degree, showing that it is at least $1 - \frac{2}{q}(\frac{1}{q} + \frac{1}{q^2} + \frac{2}{q^3})^n$ for $q \geqslant 4$. We conjecture that the density is at least $1 - \frac{1}{q}(\frac{1}{q} + \frac{1}{2q^2})^n$ for all q and n, and confirm this conjecture for dimensions $n \leqslant 37$. Finally we give a one-sided Monte Carlo algorithm called IsfCYCLIC to test whether a matrix is "good", at a cost of $\mathrm{O}(\mathrm{Mat}(n) \log n)$ field operations, where $\mathrm{Mat}(n)$ is an upper bound for the number of field operations required to multiply two matrices in $\mathrm{M}(n, \mathbb{F}_q)$.

2000 Mathematics subject classification: 15A52, 20C40

1. INTRODUCTION

The MEAT-AXE is a fundamental tool in computational representation theory, most often used to test irreducibility of a finite matrix group or algebra, and in the case of reducibility to construct an invariant subspace. A number of versions have been described in the literature, first by R. Parker [20] in 1984 and later by others [9, 16, 18]. The implementations of the MEAT-AXE in the computer algebra systems GAP [12] and Magma [3] are based on the version of D. F. Holt and S. Rees in [16]. The aim of this paper is to analyse the class of matrices used by Holt and Rees in their version of S. P. Norton's irreducibility test [16, Section 2]. In the language of Holt and Rees these are matrices whose characteristic polynomials have at least one "good" irreducible factor. Following [13] we call them *f-cyclic matrices*. They are those matrices X over $\mathbb{F}$ for

which the underlying vector space, considered as an $\mathbb{F}[X]$-module, has at least one cyclic primary component (see Section 2 for a detailed definition).

Proving that the "f-cyclic irreducibility test" is a Monte Carlo algorithm requires a lower bound on the proportion of f-cyclic matrices in an irreducible subalgebra of the algebra $\mathrm{M}(n,q)$ of $n \times n$ matrices over a field of order q. Holt and Rees derive a lower bound sufficient for their purposes by showing that at least a non-zero constant fraction of the matrices in such irreducible subalgebras have a "good" linear factor, (see [16, pp. 7-8] where a lower bound of 0.234 is proved for all n and q).

A variant of this irreducibility test using cyclic matrices was introduced by P. M. Neumann and the second author in [18], and analysing it required a lower bound for the proportion of cyclic matrices in irreducible subalgebras of $\mathrm{M}(n,q)$. Explicit lower bounds were obtained of the form $1 - cq^{-3}$ for the full matrix algebra $\mathrm{M}(n,q)$, and similar expressions for proper irreducible subalgebras, see [17, Theorems 4.1 and 5.5]. Precise limiting proportions for large n are also known, see [7, 8, 23].

In 2006 the first author began a study of f-cyclic matrices, which included both a simplified proof of the f-cyclic irreducibility test and also a determination of the exact proportion of f-cyclic matrices in $\mathrm{M}(n,q)$ for small n. The results for small n suggested that the proportion of f-cyclic matrices in $\mathrm{M}(n,q)$ may admit a lower bound $1 - cq^{-d(n)}$ for some constant c, where $d(n)$ increases with n. That is, the proportion of "non-f-cyclic" matrices may be significantly smaller than the proportion of non-cyclic matrices. Our wish to understand how this proportion varies as n increases motivated the present investigation. While the proportion of non-cyclic matrices in $\mathrm{M}(n,q)$ is known to lie between $\frac{1}{q^2(q+1)}$ and $\frac{1}{(q^2-1)(q-1)}$ for all $n \geqslant 2$ by [17, Theorem 4.1], it turns out that the proportion of non-f-cyclic matrices in $\mathrm{M}(n,q)$ decays to zero exponentially as n increases.

Theorem 1. *There is a positive constant $c < 1$ such that, for all finite field sizes q, and all dimensions $n \geqslant 1$, the proportion of f-cyclic matrices in $\mathrm{M}(n,q)$ is at least $1 - c^n$.*

It follows from our proofs that the constant $c = 0.983$ suffices for all q. Theorem 1 is proved with $c = c(q) = \mathrm{O}(q^{-1})$. We study the class of matrices that are not f-cyclic, that is to say, matrices $X \in \mathrm{M}(n,q)$ for which every primary component of the underlying vector space $\mathbb{F}_q^n$, considered as an $\mathbb{F}_q[X]$-module, is non-cyclic. We say that such matrices are *uncyclic*, and we denote by $\mathrm{unc}(n,q)$ the number of uncyclic matrices in $\mathrm{M}(n,q)$. A more precise version of our bounds follows.

Theorem 2. *If $n \geqslant 3$ and $q \geqslant 4$, then*

$$q^{-n-1}\left(1 + \left(\frac{n-1}{2}\right)q^{-1} - q^{-3}\right) < \frac{\mathrm{unc}(n,q)}{q^{n^2}} < 2q^{-1}\left(q^{-1} + q^{-2} + 2q^{-3}\right)^n.$$

The lower bound holds when $q = 2, 3$, and the following upper bounds hold

$$\frac{\mathrm{unc}(n,2)}{2^{n^2}} < (0.915)(0.983)^n \quad \textit{and} \quad \frac{\mathrm{unc}(n,3)}{3^{n^2}} < (0.52)(0.53)^n.$$

The upper bounds for this theorem are proved using induction on n, see Theorems 14 and 16. Theorem 14 involves a slightly smaller, but more elaborate, function $c^*(q)$ in place of the constant 2, see Lemma 12. Our proof of the lower bound in Theorem 2 is constructive and works for all q, see Theorem 9. We believe that the true value of $\mathrm{unc}(n,q)/q^{n^2}$ is closer to the lower bound than the upper bound given in Theorem 2, and we make the following conjecture.

Conjecture 3. *If $q \geqslant 2$ and $n \geqslant 1$, then* $\dfrac{\mathrm{unc}(n,q)}{q^{n^2}} \leqslant \dfrac{1}{q}\left(\dfrac{1}{q}+\dfrac{1}{2q^2}\right)^n$.

A different approach to estimating $\mathrm{unc}(n,q)$ is to study a probabilistic generating function for these quantities, for fixed q. We introduce such a generating function in Section 3, obtain an infinite product expansion for it in Proposition 5, and use it to compute the exact values of $\mathrm{unc}(n,q)$ as polynomials in q, for small n. These expressions are given in Table 3 for $n \leqslant 7$, and are listed in an electronic database for $n \leqslant 37$, see [15, Appendix 1]. This approach enables us to verify Conjecture 3 for $1 \leqslant n \leqslant 37$, see Proposition 8 and [15, Appendix 2].

These, to us, surprising results raise the question of whether the improved bounds for the proportion of f-cyclic matrices might lead to improvements in the MEAT-AXE algorithm. This is a matter of ongoing work of the authors, see [14]. We have resolved the first issue of whether the property of f-cyclicity can be identified efficiently. In Section 7 we give a Monte Carlo algorithm that tests whether a given matrix X in $\mathrm{M}(n,q)$ is f-cyclic, and if so constructs a generator of (possibly a direct sum of) cyclic primary summands of the underlying space considered as an $\mathbb{F}_q[X]$-module. The algorithm requires $\mathrm{O}(\mathrm{Mat}(n)\log n)$ field operations, where $\mathrm{Mat}(n)$ is an upper bound for the number of field operations required to multiply two matrices in $\mathrm{M}(n,q)$, and the construction of a constant number (depending on the desired failure probability) of random vectors in $\mathbb{F}_q^n$. For a precise statement see Theorem 18.

Section 2 gives a (known) formula for the size $|X^{\mathrm{GL}(n,q)}|$ of the $\mathrm{GL}(n,q)$-orbit containing $X \in \mathrm{M}(n,q)$ (with $\mathrm{GL}(n,q)$ acting by conjugation). The formula depends on the Frobenius canonical form of X which, in turn, depends on certain partitions. We define notation, and introduce an invariant of the $\mathrm{GL}(n,q)$-orbit called the *type* of X. In Section 3 the generating function $\sum_{n\geqslant 0}\frac{\mathrm{unc}(n,q)}{|\mathrm{GL}(n,q)|}u^n$ is expressed as an infinite product. The infinite product gives rise to a formula for $\mathrm{unc}(n,q)$ involving sums over certain partitions of rational functions in q. It not obvious from the formula that $\mathrm{unc}(n,q)$ is a polynomial in q with integer coefficients. Although the formula is explicit, we were unable to use it to prove upper bounds or lower bounds for $\mathrm{unc}(n,q)$. In Section 4 we show that $\mathrm{unc}(n,q)$ is at least $q^{n^2-n-1}+\frac{n}{2}q^{n^2-n-2}+\mathrm{O}(q^{n^2-n-3})$ by counting the number of matrices in certain large classes of uncyclic matrices. Finding upper bounds in Section 5 (for $q>2$) and in Section 6 (for $q=2$) involved a rather sensitive mathematical induction. The final Section 7 gives a practical Monte Carlo $\mathrm{O}(\mathrm{Mat}(n)\log n)$ algorithm to test whether a given matrix X is f-cyclic relative to some irreducible divisor of $c_X(t)$. This algorithm avoids

the expensive step of evaluating a divisor of $c_X(t)$ at X. Moreover, it outputs a (witness) vector u which can be used when applying Norton's irreducibility test [16, Section 2.1].

2. Conjugacy Classes in $\mathrm{GL}(n,q)$

A partition of $n \in \mathbb{N} := \{0, 1, 2, \dots\}$, written $\lambda \vdash n$, is an unordered sum $n = \sum_{i \geqslant 1} \lambda_i$ where the parts λ_i lie in $\mathbb{N}$. A partition can be represented by (a) its parts, (b) its Young (or Ferrers) diagram [21], or (c) by the multiplicities of its parts. We write $\lambda = (\lambda_1, \lambda_2, \dots)$ where $\lambda_1 \geqslant \lambda_2 \geqslant \cdots$ and $n = \sum_{i \geqslant 1} \lambda_i$. Set $|\lambda| := \sum_{i \geqslant 1} \lambda_i$. It is convenient to abbreviate a partition by omitting all (or some) of the trailing zeroes. We shall commonly write $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_k)$ where $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \lambda_k > 0$ and $\lambda_{k+1} = \lambda_{k+2} = \cdots = 0$. The empty partition, or partition of zero, is written $(0, 0, \dots)$ or simply $()$.

The Young diagram of $\lambda = (\lambda_1, \dots, \lambda_k)$ is a rectangular array of $|\lambda|$ boxes arranged in k left-justified rows, with λ_i boxes in row i, for each i. For example, Figure 1 shows the Young diagrams for the partitions $\lambda = (5, 3, 3, 1)$ and $\mu = (4, 3, 3, 1, 1)$ of $n = 12$.

Figure 1. Young diagrams for $\lambda = (5, 3, 3, 1)$ (left) and $\mu = (4, 3, 3, 1, 1)$ (right).

By interchanging the rows and columns of the Young diagram of λ, we obtain the Young diagram of another partition, called the *conjugate* partition, and denoted λ'. For example, in Figure 1, $\lambda' = \mu$ and $\mu' = \lambda$. The number of parts of λ equal to i, that is to say, the multiplicity of i, is denoted $m_i(\lambda)$ or simply m_i. We occasionally write $\lambda = 1^{m_1} 2^{m_2} 3^{m_3} \cdots$. The number of non-zero parts of λ, written $\ell_1(\lambda)$, is the number of squares in the first column of the Young diagram of λ. More generally, $\ell_i(\lambda)$ denotes the number of squares in the first i columns of the Young diagram of λ.

The vector $m(\lambda) = (m_1(\lambda), m_2(\lambda), \dots)$ in $\mathbb{N}^\infty$ need not be a partition because the coordinates need not satisfy $m_i(\lambda) \geqslant m_{i+1}(\lambda)$ for $i \geqslant 1$. Denote by $m(\lambda)^{\mathrm{o}}$ the partition obtained from $m(\lambda)$ by permuting the coordinates so that they are weakly decreasing. The formula for the order $|\mathrm{C}_{\mathrm{GL}(n,q)}(X)|$ of the centralizer of an element $X \in \mathrm{M}(n,q)$ involves three vectors: $m(\lambda)$, $\ell(\lambda) := (\ell_1(\lambda), \ell_2(\lambda), \dots)$, and $e(\lambda) := (m(\lambda)^{\mathrm{o}})'$, for various partitions λ, see (1) and (2) below. As an example, if $\lambda = (5, 3, 3, 1)$, then $\lambda' = (4, 3, 3, 1, 1)$, and

$$m(\lambda) = (1, 0, 2, 0, 1, 0, \dots), \quad \ell(\lambda) = (4, 7, 10, 11, 12, 12, \dots), \quad e(\lambda) = (3, 1, 0, \dots).$$

The reader should not confuse the vector $e(\lambda)$ with the symmetric polynomial e_λ defined in [21, p. 290]. It is convenient to define the dot product $x \cdot y := \sum_{i \geqslant 1} x_i y_i$ of vectors $x, y \in \mathbb{N}^\infty$ in the case that the sum is finite, for example, when x or y has finite support. Also define $\| x \|^2 := x \cdot x$.

Lemma 4. *Let λ be a partition of $|\lambda|$. Then*

(a) $|\lambda| = \sum_{i\geqslant 1} i\, m_i(\lambda) = |\lambda'|$,
(b) $m_i(\lambda) = \lambda'_i - \lambda'_{i+1}$,
(c) $\ell_i(\lambda) = \lambda'_1 + \cdots + \lambda'_i = \left(\sum_{k<i} k m_k(\lambda)\right) + i\left(\sum_{k\geqslant i} m_k(\lambda)\right)$,
(d) $m(\lambda)\cdot\ell(\lambda) = \| \lambda' \|^2 \equiv |\lambda| \pmod 2$,
(e) $e_k(\lambda) = |\{i \mid m_i(\lambda) \geqslant k\}|$,
(f) $\| \lambda' \|^2 \geqslant |\lambda|$ *with equality if and only if* $\lambda = (|\lambda|, 0, 0, \dots)$,
(g) $|e(\lambda)| = \sum_{i\geqslant 1} m_i(\lambda) = \lambda'_1$.

Proof. The proofs of parts (a),(b) are elementary, see [21, p. 287]. Counting the squares in the first i columns of the Young diagram for λ by columns gives the first formula for $\ell_i(\lambda)$ in part (c), while counting by rows gives the second. Consider part (d):

$$\begin{aligned} m(\lambda)\cdot\ell(\lambda) &= \sum_{i\geqslant 1} m_i(\lambda)\ell_i(\lambda) \\ &= (\lambda'_1 - \lambda'_2)\ell_1(\lambda) + \sum_{i\geqslant 2}(\lambda'_i - \lambda'_{i+1})\ell_i(\lambda) && \text{by part (b)} \\ &= (\lambda'_1 - \lambda'_2)\lambda'_1 + \sum_{i\geqslant 2}\left(\lambda'_i\ell_{i-1}(\lambda) + (\lambda'_i)^2 - \lambda'_{i+1}\ell_i(\lambda)\right) && \text{as } \ell_i(\lambda) = \ell_{i-1}(\lambda) + \lambda'_i \\ &= -\lambda'_2\lambda'_1 + \lambda'_2\lambda'_1 + \sum_{i\geqslant 1}(\lambda'_i)^2 = \| \lambda' \|^2 . \end{aligned}$$

However, $(\lambda'_i)^2 \equiv \lambda'_i \pmod 2$ and so $\| \lambda' \|^2 \equiv \sum_{i\geqslant 1}\lambda'_i \pmod 2$. Part (d) now follows as $|\lambda'| = |\lambda|$. Part (e) follows from the elementary fact $\lambda'_k = |\{i \mid \lambda_i \geqslant k\}|$, while part (f) follows from (d) and the observation that $\lambda_i'^2 \geqslant \lambda'_i$ with equality if and only if $\lambda'_i = 0, 1$. Finally, part (g) follows as $\sum_{i\geqslant 1} m_i(\lambda)$ and λ'_1 both count the number of rows in the Young diagram of λ, and $e(\lambda) = (m(\lambda)^{\circ})'$ so $\sum_{i\geqslant 1} e_i(\lambda) = \sum_{i\geqslant 1} m_i(\lambda)$. □

Recall that $\mathrm{M}(n, q)$ is the algebra of $n \times n$ matrices over $\mathbb{F}_q$, and let $G = \mathrm{GL}(n, q)$ denote the general linear group, its group of units. A formula for the size $|X^G|$ of the G-orbit of a matrix $X \in \mathrm{M}(n, q)$ dates back at least to [11, 22]. Our formula is better suited for calculation. Clearly, $|X^G| = |G : C_G(X)|$ and the structure of the centralizer $C_G(X)$ of X depends on the Frobenius (or rational) canonical form of X. Suppose that the characteristic polynomial $c_X(t)$ factors as $\prod_f f^{\nu(f)}$ where the product is over monic irreducible polynomials $f(t) \in \mathbb{F}_q[t]$, and $\nu(f) \in \mathbb{N}$ (possibly $\nu(f) = 0$). The structure of $C_G(X)$ depends on partitions $\lambda(f, X)$ of $\nu(f)$ which we abbreviate $\lambda(f)$ when the dependence on X is clear, see [11, 22]. The vector space $V = \mathbb{F}_q^{1\times n}$ is an $\mathbb{F}_q[X]$-module, and $V(f) = \ker f^{\nu(f)}(X) = \ker f(X)^{\lambda(f)_1}$ is its f-primary component. Let $X(f)$ denote the restriction of X to $V(f)$. Thus the minimal polynomial of $X(f)$ is $f^{\lambda(f)_1}$, and that of X is $m_X(t) = \prod_f f^{\lambda(f)_1}$. Now X is conjugate to a block diagonal matrix $\bigoplus X(f)$ and

$V(f)$ is isomorphic as an $\mathbb{F}_q[X]$-module to

$$V(f) \cong \bigoplus_{i \geqslant 1} \mathbb{F}_q[t]/(f(t)^{\lambda(f)_i}).$$

Two matrices X and Y lie in the same G-orbit if and only if they have the same Frobenius canonical form, that is, if and only if $\lambda(f, X) = \lambda(f, Y)$ for all monic irreducibles f. It is convenient to define a formal expression called the *type* of X written $\text{type}(X) := \prod_f f^{\lambda(f,X)}$. Two formal expressions of this kind are regarded as equal if and only if their respective exponent partitions are equal. Thus X and Y lie in the same G-orbit if and only if $\text{type}(X) = \text{type}(Y)$. As it is sometimes convenient to omit trivial factors f^0 from the product $c_X(t) = \prod_f f^{|\lambda(f,X)|}$, it is therefore sometimes convenient to omit factors $f^{(0,0,\dots)}$ from $\text{type}(X)$.

It follows from [11, 22] that

$$|\mathrm{C}_{\mathrm{GL}(n,q)}(X)| = \prod_f |\mathrm{C}_{\mathrm{GL}(V(f))}(X(f))| = \prod_f c(\lambda(f), q^{d(f)}) \tag{1}$$

where $d(f) := \deg(f)$ and $c(\lambda, q)$ is the function

$$c(\lambda, q) := \prod_{i=1}^{\lambda_1} \prod_{k=1}^{m_i(\lambda)} (q^{\ell_i(\lambda)} - q^{\ell_i(\lambda)-k}) = q^{m(\lambda)\cdot\ell(\lambda)} \prod_{i=1}^{\lambda_1} \prod_{k=1}^{m_i(\lambda)} (1 - q^{-k}),$$

see [11, 22]. By Lemma 4(d) and (e), $c(\lambda, q)$ may be rewritten as

$$c(\lambda, q) = q^{\|\lambda'\|^2} \prod_{k \geqslant 1} (1 - q^{-k})^{e_k(\lambda)}. \tag{2}$$

In summary,

$$|X^{\mathrm{GL}(n,q)}| = |\mathrm{GL}(n, q)| \prod_f \frac{1}{c(\lambda(f), q^{d(f)})}. \tag{3}$$

The following table of values of $c(\lambda, q)$ both illustrates formula (2), and provides data for the proof of Lemma 11. In this table we shall assume $\lambda_1 > \lambda_2 > \lambda_3$, and we use the notation $1^{m_1}2^{m_2}\cdots$ to indicate multiplicities $m(\lambda) = (m_1, m_2, \dots)$. For example, (λ_1, λ_2) is written as $\lambda_1^1\lambda_2^1$ because λ_1 and λ_2 each occur once, given our assumption $\lambda_1 > \lambda_2$.

For a monic irreducible polynomial g over $\mathbb{F}$, a matrix $X \in \mathrm{M}(n, \mathbb{F})$ is said to be *f-cyclic relative to g* if the restriction $X(g)$ of X to the g-primary component $V(g)$ of $V = \mathbb{F}^{1\times n}$ is cyclic. Although we are interested to count matrices X that are f-cyclic relative to some monic irreducible divisor g of $c_X(t)$, the complementary count is easier. We call X *uncyclic* if $X(g)$ is not cyclic for *all* monic irreducible divisors g of $c_X(t)$. Equivalently, X is uncyclic if and only if $\lambda(g)'_1 \neq 1$ for all g (that is, $\lambda(g)$ has zero or at least two parts for each g). One can readily see from the factorizations $c_X(t) = \prod g^{\nu(g)}$ and $m_X(t) = \prod g^{\mu(g)}$ of the characteristic and minimal polynomials of X whether or not X is f-cyclic (or

λ	$\lvert\lambda\rvert$	λ'	$e(\lambda)$	$c(\lambda, q)$
$\lambda_1^1\lambda_2^1$	$\lambda_1+\lambda_2$	$1^{\lambda_1-\lambda_2}2^{\lambda_2}$	(2)	$q^{\lvert\lambda\rvert+2\lambda_2}(1-q^{-1})^2$
λ_1^2	$2\lambda_1$	2^{λ_1}	$(1,1)$	$q^{2\lvert\lambda\rvert}(1-q^{-1})(1-q^{-2})$
$\lambda_1^1\lambda_2^1\lambda_3^1$	$\lambda_1+\lambda_2+\lambda_3$	$1^{\lambda_1-\lambda_2}2^{\lambda_2-\lambda_3}3^{\lambda_3}$	(3)	$q^{\lvert\lambda\rvert+2\lambda_2+4\lambda_3}(1-q^{-1})^3$
$\lambda_1^1\lambda_2^2$	$\lambda_1+2\lambda_2$	$1^{\lambda_1-\lambda_2}3^{\lambda_2}$	$(2,1)$	$q^{\lvert\lambda\rvert+6\lambda_2}(1-q^{-1})^2(1-q^{-2})$
$\lambda_1^2\lambda_2^1$	$2\lambda_1+\lambda_2$	$2^{\lambda_1-\lambda_2}3^{\lambda_2}$	$(2,1)$	$q^{2\lvert\lambda\rvert+3\lambda_2}(1-q^{-1})^2(1-q^{-2})$
λ_1^k	$k\lambda_1$	k^{λ_1}	1^k	$q^{\lambda_1 k^2}\prod_{i=1}^{k}(1-q^{-i})$

TABLE 1. Values of $c(\lambda, q)$.

uncyclic): f-cyclic relative to g means $\nu(g) = \mu(g)$, and uncyclic means that, for all g, $\nu(g) > \mu(g) > 0$ or $\nu(g) = \mu(g) = 0$.

3. Generating function as an infinite product

In this section we express the generating function

$$\mathrm{Unc}_q(u) := 1 + \sum_{n=1}^{\infty} \frac{\mathrm{unc}(n,q)}{|\mathrm{GL}(n,q)|}\, u^n \tag{4}$$

as an infinite product. It is more convenient to consider the weighted proportion $\frac{\mathrm{unc}(n,q)}{|\mathrm{GL}(n,q)|}$ of uncyclic matrices in $\mathrm{M}(n,q)$ because orbit sizes have a factor $|\mathrm{GL}(n,q)|$ in the numerator.

Our main tool is the cycle index for $\mathrm{M}(n,q)$ which is defined as

$$Z_{\mathrm{M}(n,q)} := \frac{1}{|\mathrm{GL}(n,q)|} \sum_{X \in \mathrm{M}(n,q)} \left(\prod_f x_{f,\lambda(X,f)} \right) \tag{5}$$

where the product is over all monic irreducible polynomials and the $x_{f,\lambda}$ are indeterminates, see [11, 22] and [6, pp. 35-36]. If we set $x_{f,()} := 1$ for each f, then for each X the product in (5) has finitely many factors different to 1.

Stong [22], building on the work of Kung [11], proves that

$$1 + \sum_{n=1}^{\infty} Z_{\mathrm{M}(n,q)} u^n = \prod_f \left(\sum_\lambda x_{f,\lambda} \frac{u^{|\lambda| d(f)}}{c(\lambda, q^{d(f)})} \right) \tag{6}$$

where the sum on the right-hand side is over all partitions $(), (1), (2), (1,1), \ldots$. By convention

$$c((), q) = |\mathrm{GL}(0,q)| = \mathrm{unc}(0,q) = 1.$$

Proposition 5. *Let Λ_1 be the set of partitions λ such that $\lambda'_1 \neq 1$ (equivalently λ has 0 or at least 2 parts). Then*

$$\mathrm{Unc}_q(u) = \sum_{n\geqslant 0} \frac{\mathrm{unc}(n,q)}{|\mathrm{GL}(n,q)|} u^n = \prod_f \left(\sum_{\lambda\in\Lambda_1} \frac{u^{|\lambda(f)|d(f)}}{c(\lambda(f), q^{d(f)})} \right). \tag{7}$$

Proof. From the remarks above, X is uncyclic if and only if $\lambda(f) \in \Lambda_1$ for all f. As the set of uncyclic matrices in $\mathrm{M}(n,q)$ is a union of $\mathrm{GL}(n,q)$-orbits, it follows from (3) that

$$\mathrm{unc}(n,q) = \sum |\mathrm{GL}(n,q)| \prod_f \frac{1}{c(\lambda(f), q^{d(f)})}$$

where the sum ranges over all decompositions $n = \sum |\lambda(f)|d(f)$ with $\lambda(f) \in \Lambda_1$. This proves (7).

An alternative proof uses (6). In (5) set $x_{f,\lambda} = 1$ if $\lambda \in \Lambda_1$, and 0 otherwise. Then $Z_{\mathrm{M}(n,q)}$ equals $\mathrm{unc}(n,q)/|\mathrm{GL}(n,q)|$. On the other hand, the bracketed sums of (6) and (7) are equal. □

As the bracketed sum in (7) is the same for all f with degree r, we define

$$A(q,u) := \sum_{\lambda\in\Lambda_1} \frac{u^{|\lambda|}}{c(\lambda,q)} \quad \text{and} \quad a_n(q) := \sum_{\lambda\vdash n,\, \lambda\neq(n)} \frac{1}{c(\lambda,q)}. \tag{8}$$

Thus $A(q,u) = \sum_{n\geqslant 0} a_n(q)u^n$ where $a_0(q) = 1$, $a_1(q) = 0$, $a_2(q) = |\mathrm{GL}(2,q)|^{-1}$, etc. Denote by $N(r,q)$ the number of monic irreducible polynomials over $\mathbb{F}_q$ of degree r. Then (7) may be rewritten

$$\mathrm{Unc}_q(u) = \sum_{n\geqslant 0} \frac{\mathrm{unc}(n,q)}{|\mathrm{GL}(n,q)|} u^n = \prod_{r\geqslant 1} A(q^r, u^r)^{N(r,q)} = \prod_{r\geqslant 1} \left(1 + \sum_{n\geqslant 2} a_n(q^r) u^{rn} \right)^{N(r,q)}. \tag{9}$$

A closed formula for $\mathrm{unc}(n,q)$ can be obtained by expanding the products in (9). This formula, though unwieldy, may be used to to determine $\mathrm{unc}(n,q)$ for small n.

Lemma 6. *Given $n \in \mathbb{N}$ and a partition $\lambda = 1^{m_1}2^{m_2}\cdots$ with $\lambda'_1 \leqslant n$, denote the multinomial coefficient $\binom{n}{n-\sum_{i\geqslant 1} m_i, m_1, m_2, \ldots} = \frac{n!}{(n-\lambda'_1)!m_1!m_2!\cdots}$ by $\binom{n}{m(\lambda)}$. Then*

$$(1 + a_1u + a_2u^2 + \cdots)^n = \sum_{k\geqslant 0} \left(\sum_{\lambda\vdash k} \binom{n}{m(\lambda)} a^{m(\lambda)} \right) u^k \tag{10}$$

where $a^{m(\lambda)} := a_1^{m_1}a_2^{m_2}\cdots$.

Proof. Set $a_0 := 1$. Expanding the left-hand side of (10) gives

$$\sum_{\lambda\in\mathbb{N}^n} a_{\lambda_1}u^{\lambda_1}a_{\lambda_2}u^{\lambda_2}\cdots a_{\lambda_n}u^{\lambda_n} = \sum_{k\geqslant 0}\left(\sum_{\lambda\in\mathbb{N}^n,\,|\lambda|=k} a_{\lambda_1}a_{\lambda_2}\cdots a_{\lambda_n}\right)u^k. \tag{11}$$

The term $a_{\lambda_1}\cdots a_{\lambda_n}$ will be repeated $\binom{n}{m(\lambda)}$ times, where $\binom{n}{m(\lambda)}$ is the number of distinct elements of $\mathbb{N}^n$ obtained by permuting the coordinates of $\lambda = (\lambda_1,\ldots,\lambda_n)$. If $1^{m_1}2^{m_2}\cdots$ is the unique partition corresponding to λ, then $a_{\lambda_1}\cdots a_{\lambda_n} = a^{m(\lambda)}$ because a_i has multiplicity m_i for $i \geqslant 1$, and multiplicity $n - \sum_{i\geqslant 1} m_i = n - \lambda'_1$ for $i = 0$ by Lemma 4(g). □

Lemma 6 may be used to expand the powers in (9). Since in (9) we have $a_1 = 0$, it follows from (11) that the inner sum in (10) is over partitions λ of k with no part of size 1. For example, if $k = 5$, then $\lambda = (5)$ or $(3,2)$ and $\binom{n}{m(\lambda)}$ equals n or $n(n-1)$, respectively. Expanding the power $(1 + a_2z^2 + a_3z^3 + \cdots)^n$ using Lemma 6 gives

$$\begin{aligned}
&1 + na_2z^2 + na_3z^3 + \left(na_4 + \binom{n}{2}a_2^2\right)z^4 + \left(na_5 + 2\binom{n}{2}a_2a_3\right)z^5 + \cdots\\
&= 1 + n\left(\sum_i a_iz^i\right) + \binom{n}{2}\left(\sum_i a_i^2z^{2i} + 2\sum_{i<j} a_ia_jz^{i+j}\right)\\
&\quad + \binom{n}{3}\left(\sum_i a_i^3z^{3i} + 3\sum_{i<j} a_i^2a_jz^{2i+j} + 3\sum_{i<j} a_ia_j^2z^{i+2j} + 6\sum_{i<j<k} a_ia_ja_kz^{i+j+k}\right) + \cdots.
\end{aligned}$$

In order to evaluate (9) it is useful to substitute $z = u^r$ and $n = N(r,q)$ in the above expression. By using (10) and (9) one can, in principle, write down a closed form for $\mathrm{unc}(n,q)$. The resulting closed form is rather complicated, and it is not obviously useful for bounding $\mathrm{unc}(n,q)$. In [15, Appendix 2] we give a MAGMA [3] computer program for computing $\mathrm{unc}(n,q)$ for small n. Given that the number of partitions of n (even those with no part of size 1) is asymptotically exponential (see [2, p. 70]), our computer program can compute $\mathrm{unc}(n,q)$ only for small n.

For very small values of n one does not need a computer program. Equating the coefficient of u^n for $n \leqslant 5$ on both sides of (9) gives values of $\frac{\mathrm{unc}(n,q)}{|\mathrm{GL}(n,q)|}$ in terms of the polynomials $a_n(q)$ defined in (8). This information is summarized in Table 2.

It is easy to show that $\mathrm{unc}(1,q) = 0$. The values of $\mathrm{unc}(n,q)$ for $n = 2,3,4,5$ can be computed from Table 2. We list the values and $\mathrm{unc}(n,q)$ for $n \leqslant 7$ in Table 3 below.

The polynomials $\mathrm{unc}(n,q)$ for $n \leqslant 37$ were computed with the MAGMA [3] programs in [15, Appendix 2] and stored in the database [15, Appendix 1]. Lemma 7 below is useful for bounding polynomials in q (or q^{-1}).

n	$\frac{\mathrm{unc}(n,q)}{\lvert\mathrm{GL}(n,q)\rvert}$	$a_n(q)$
2	$\binom{q}{1}a_2(q)$	$\frac{1}{c((1,1),q)} = \frac{1}{q^4-q^3-q^2+q}$
3	$\binom{q}{1}a_3(q)$	$\frac{1}{c((1,1,1),q)} + \frac{1}{c((2,1),q)} = \frac{q^3+q^2-1}{q^8-q^7-q^6+q^4+q^3-q^2}$
4	$\binom{q}{1}a_4(q) + \binom{q}{2}a_2(q) + N(2,q)a_2(q^2)$	$\frac{q^7+q^6+q^5-q^4-q^3-q^2+1}{q^{13}-q^{12}-q^{11}+2q^8-q^5-q^4+q^3}$
5	$\binom{q}{1}a_5(q) + q(q-1)a_2(q)a_3(q)$	$\frac{q^{12}+q^{11}+q^{10}-q^8-2q^7-q^6+q^4+q^3+q^2-1}{q^{19}-q^{18}-q^{17}+q^{14}+q^{13}+q^{12}-q^{11}-q^{10}-q^9+q^6+q^5-q^4}$

TABLE 2. Values of $\frac{\mathrm{unc}(n,q)}{\lvert\mathrm{GL}(n,q)\rvert}$ and $a_n(q)$ for $2 \leqslant n \leqslant 5$.

n	$\mathrm{unc}(n,q)$
2	q
3	$q^5+q^4-q^2$
4	$q^{11}+2q^{10}-2q^7-q^5+q^4$
5	$q^{19}+2q^{18}+2q^{17}+q^{16}-q^{15}-2q^{14}-3q^{13}-q^{12}+q^{10}+q^9+q^8-q^7$
6	$q^{29}+3q^{28}+3q^{27}+3q^{26}-q^{25}-5q^{23}-5q^{22}-3q^{21}-2q^{20}+2q^{18}+4q^{17}+3q^{15}$ $-q^{14}-2q^{12}+q^{11}$
7	$q^{41}+3q^{40}+5q^{39}+5q^{38}+3q^{37}-4q^{35}-9q^{34}-11q^{33}-12q^{32}-7q^{31}-3q^{30}+4q^{29}$ $+6q^{28}+11q^{27}+8q^{26}+7q^{25}+q^{23}-3q^{22}-2q^{21}-3q^{20}+2q^{17}-q^{16}$

TABLE 3. Values of $\mathrm{unc}(n,q)$ for $2 \leqslant n \leqslant 7$.

Lemma 7. *Suppose that m, n are positive integers and $\alpha_0, \alpha_1, \ldots, \alpha_{m-1}$, $\beta_0, \beta_1, \ldots, \beta_{n-1}$ are non-negative real numbers. Set*

$$c(q) := (\alpha_{m-1}q^{m-1} + \cdots + \alpha_1 q + \alpha_0)q^n - (\beta_{n-1}q^{n-1} + \cdots + \beta_1 q + \beta_0).$$

If $q_0 \geqslant 0$ and $c(q_0) \geqslant 0$, then $c(q) \geqslant 0$ for all $q \geqslant q_0$.

Proof. Set $a(q) := \alpha_{m-1}q^{m-1}+\cdots+\alpha_1 q+\alpha_0$, and $b(q) := \beta_{n-1}q^{-1}+\cdots+\beta_1 q^{-(n-1)}+\beta_0 q^{-n}$. Then $c(q) = (a(q) - b(q))q^n$. Since $a(q) \geqslant a(q_0)$ and $b(q_0) \geqslant b(q)$, it follows that $a(q) - b(q) \geqslant a(q_0) - b(q_0)$ and so $c(q) \geqslant c(q_0) \geqslant 0$. Thus $c(q) \geqslant 0$ for all $q \geqslant q_0$. □

Lemma 7 may be applied to verify Conjecture 3 for small n.

Proposition 8. *If $q \geqslant 2$ and $1 \leqslant n \leqslant 37$, then $\mathrm{unc}(n,q) \leqslant q^{n^2-n-1}(1+\frac{1}{2q})^n$.*

Proof. The idea is to list the difference polynomials $d_n(q) = q^{n^2-n-1}(1+\frac{1}{2q})^n - \mathrm{unc}(n,q)$ for $1 \leqslant n \leqslant 37$ and repeatedly apply Lemma 7. For example, $d_5(q)$ equals

$$d_5(q) = \frac{1}{2}q^{18} + \frac{1}{2}q^{17} + \frac{1}{4}q^{16} + \frac{21}{16}q^{15} + \frac{65}{32}q^{14} + 3q^{13} + q^{12} - q^{10} - q^9 - q^8 + q^7,$$

and Lemma 7 with $q_0 = 2$ shows that polynomials $d_5(q) - q^7$ and q^7 are both non-negative for $q \geqslant 2$. Adding shows $d_5(q) \geqslant 0$ for $q \geqslant 2$. For more a complicated polynomial such as $q^8 - 3q^6 + q^5 - 5q^4$, Lemma 7 shows $q^8 - 3q^6 \geqslant 0$ for $q \geqslant 2$ and $q^5 - 5q^4 \geqslant 0$ for $q \geqslant 5$. Thus $q^8 - 3q^6 + q^5 - 5q^4 \geqslant 0$ holds for $q \geqslant 5$. Evaluating at $q = 2, 3, 4$ shows that $q^8 - 3q^6 + q^5 - 5q^4 \geqslant 0$ holds for $q \geqslant 2$. The MAGMA [3] computer program listed in [15, Appendix 2] uses these ideas to verify Conjecture 3 for $1 \leqslant n \leqslant 37$. □

4. A lower bound for $\mathrm{unc}(n, q)$

In this section we count the uncyclic matrices $X \in \mathrm{M}(n, q)$ with $\mathrm{type}(X) = (t-\alpha)^\lambda$ or $\mathrm{type}(X) = (t-\alpha)^\lambda(t-\beta)^\mu$, where α, β are distinct elements of $\mathbb{F}_q$, and λ, μ are partitions with $|\lambda| = n$ or $|\lambda| + |\mu| = n$ respectively (recall the definition of $\mathrm{type}(X)$ preceding (1)). If Conjecture 3 were correct, then it would follow from the binomial theorem that

$$\mathrm{unc}(n, q) \leqslant q^{n^2-n-1} + \frac{n}{2} q^{n^2-n-2} + \mathrm{O}(q^{n^2-n-3})$$

where the constant involved in $\mathrm{O}(q^{n^2-n-3})$ is independent of q. The main result of this section is that there is a lower bound comparable to this conjectured upper bound.

Theorem 9. *If $q \geqslant 2$ and $n \geqslant 3$, then*

$$q^{n^2-n-1}\left(1 + \left(\frac{n-1}{2}\right) q^{-1} - q^{-3}\right) < \mathrm{unc}(n, q).$$

The proof uses the quantity $\omega(n, q) := \prod_{i=1}^{n}(1 - q^{-i}) = q^{-n^2}|\mathrm{GL}(n, q)|$.

Lemma 10. *If $n \geqslant 1$, then $(1 - q^{-1})^2 < 1 - q^{-1} - q^{-2} < \omega(\infty, q) < \omega(n, q) \leqslant 1 - q^{-1}$.*

Proof. See Lemma 3.5 and Corollary 3.6 of [17]. □

Let $\alpha \in \mathbb{F}_q$. A matrix $X \in \mathrm{M}(n, q)$ is *α-potent* if its characteristic polynomial is $c_X(t) = (t - \alpha)^n$. The map $X \to X + (\beta - \alpha)I$ is a bijection between the subsets of α-potent matrices and β-potent matrices in $\mathrm{M}(n, q)$. In particular, the numbers of α-potent and unipotent matrices in $\mathrm{M}(n, q)$ are equal. The number of unipotent matrices in $\mathrm{M}(n, q)$ (or in $\mathrm{GL}(n, q)$) equals $q^{n(n-1)}$ by a theorem of Steinberg [4, Theorem 6.6.1]. Denote by $\mathrm{U}(n, q, \alpha)$ the *set of uncyclic α-potent matrices* in $\mathrm{M}(n, q)$. Note that $X \in \mathrm{U}(n, q, \alpha)$ if and only if $\mathrm{type}(X) = (t - \alpha)^\lambda$ where λ has more than one part.

Let $r(n, q)$ denote the number of uncyclic matrices X in $\mathrm{M}(n, q)$ with $\mathrm{type}(X) = f^\lambda$ where f is a monic irreducible polynomial whose degree divides n. Let $r(n, q, d)$ denote the number of such matrices X where $\mathrm{type}(X) = f^\lambda$, and f has degree d for a fixed divisor d of n. Thus $r(n, q) = \sum_{d|n} r(n, q, d)$. Estimating the size of $r(n, q, 1)$, is an important step towards estimating $r(n, q)$, which, in turn, will help us bound $\mathrm{unc}(n, q)$.

Lemma 11. *Let $r(n,q,1)$ denote the number of uncyclic matrices in $\mathrm{M}(n,q)$ that are α-potent for some $\alpha \in \mathbb{F}_q$. If $n \geqslant 1$, then $r(n,q,1) = c_0(n,q)q^{n^2-n-1}$ where*

$$c_0(n,q) := q^2\left(1 - \prod_{i=2}^{n}(1-q^{-i})\right).$$

Moreover, $1 + q^{-1} - q^{-3} \leqslant c_0(n,q) < 1 + q^{-1} + q^{-2}$ for $n \geqslant 3$, and $\lim_{q\to\infty} c_0(n,q) = 1$.

Proof. Since $|\mathrm{U}(n,q,\alpha)|$ is independent of $\alpha \in \mathbb{F}_q$, it follows that $r(n,q,1) = q|\mathrm{U}(n,q,1)|$. Thus it remains to count the uncyclic unipotent matrices. A cyclic unipotent matrix belongs to a conjugacy class with type $(t-1)^{(n)}$, and an uncyclic unipotent matrix X has $\mathrm{type}(X) = (t-1)^\lambda$ for some $\lambda \neq (n)$. By (2), the centralizer of a cyclic unipotent matrix has order $q^n(1-q^{-1})$. It follows, using the above mentioned theorem of Steinberg, that

$$|\mathrm{U}(n,q,1)| = q^{n(n-1)} - \frac{q^{n^2}\prod_{i=1}^{n}(1-q^{-i})}{q^n(1-q^{-1})} = q^{n^2-n}\left(1 - \prod_{i=2}^{n}(1-q^{-i})\right).$$

The cardinality of the disjoint union $\bigcup_{\alpha\in\mathbb{F}_q}\mathrm{U}(n,q,\alpha)$ is thus

$$r(n,q,1) = q^{n^2-n+1}\left(1 - \prod_{i=2}^{n}(1-q^{-i})\right) = c_0(n,q)q^{n^2-n-1}.$$

It remains to estimate $c_0(n,q)$. Since $c_0(n,q)$ is an increasing function of n, it follows that $1 + q^{-1} - q^{-3} = c_0(3,q) \leqslant c_0(n,q) < c_0(\infty,q)$ for $n \geqslant 3$. The following calculation shows that the limit

$$c_0(\infty,q) = 1 + q^{-1} + q^{-2} - q^{-5} - q^{-6} - q^{-7} - q^{-8} - q^{-9} + q^{-13} + q^{-14} + \cdots \tag{12}$$

is finite for all q:

$$\begin{aligned} c_0(\infty,q) &= q^2\left[1 - (1-q^{-2})\prod_{i\geqslant 3}(1-q^{-i})\right] < q^2\left[1 - (1-q^{-2})\left(1 - \sum_{i\geqslant 3} q^{-i}\right)\right] \\ &= q^2\left[1 - (1-q^{-2})\left(1 - \frac{q^{-3}}{1-q^{-1}}\right)\right] = 1 + q^{-1} + q^{-2}. \end{aligned}$$

Finally, $\lim_{q\to\infty}(1 + q^{-1} - q^{-3}) = \lim_{q\to\infty}(1 + q^{-1} + q^{-2}) = 1$ so $\lim_{q\to\infty} c_0(n,q) = 1$. □

Proof of Theorem 9. By Lemma 11 the number $r(n,q,1)$ of uncyclic matrices in $\mathrm{M}(n,q)$ with type $(t-\alpha)^\lambda$, for some $\alpha \in \mathbb{F}_q$ and $\lambda \neq (n)$, is $q^{n^2-n-1} + q^{n^2-n-2} + \mathrm{O}(q^{n^2-n-3})$. We shall now show that the number of uncyclic matrices in $\mathrm{M}(n,q)$ with type $(t-\alpha)^\lambda(t-\beta)^\mu$ where $\alpha \neq \beta$ is $\left(\frac{n-3}{2}\right)q^{n^2-n-2} + \mathrm{O}(q^{n^2-n-3})$. These two contributions give a lower bound for $\mathrm{unc}(n,q)$ approximately of the size forecast in the preamble to this section.

It is easy to check using the values for $\mathrm{unc}(n,q)$ in Table 3 that Theorem 9 is true for $n = 3, 4$. Assume henceforth that $n \geqslant 5$. We count the number of matrices $X \in \mathrm{M}(n,q)$

with type$(X) = (t-\alpha)^{(\lambda_1,1)}(t-\beta)^{(\mu_1,1)}$, for fixed elements $\alpha \neq \beta$ in $\mathbb{F}_q$ and $\lambda_1 \geqslant \mu_1 \geqslant 1$ such that $n = \lambda_1 + \mu_1 + 2$. It follows from Table 1 that

$$c((\lambda_1,1),q) = \begin{cases} q^{\lambda_1+3}(1-q^{-1})^2 & \text{if } \lambda_1 > 1, \\ q^{\lambda_1+3}(1-q^{-1})(1-q^{-2}) & \text{if } \lambda_1 = 1. \end{cases}$$

Since $\lambda_1 + 3 + \mu_1 + 3 = n + 4$, it follows from (3) that X lies in a $\mathrm{GL}(n,q)$-orbit of size

$$\frac{q^{n^2-n-4}\omega(n,q)}{(1-q^{-1})^4}, \qquad \frac{q^{n^2-n-4}\omega(n,q)}{(1-q^{-1})^3(1-q^{-2})}, \quad \text{or} \quad \frac{q^{n^2-n-4}\omega(n,q)}{(1-q^{-1})^2(1-q^{-2})^2}$$

if $\lambda_1 \geqslant \mu_1 > 1$, $\lambda_1 > \mu_1 = 1$, or $\lambda_1 = \mu_1 = 1$, respectively.

How many $\mathrm{GL}(n,q)$-orbits arise if we vary $\alpha \neq \beta$ and $\lambda_1 \geqslant \mu_1 \geqslant 1$? To answer this question we consider three cases: (a) $\lambda_1 > \mu_1 > 1$, (b) $\lambda_1 = \mu_1 > 1$, and (c) $\lambda_1 > \mu_1 = 1$. (As $n \geqslant 5$, the case $\lambda_1 = \mu_1 = 1$ does not arise.) (a) If $\lambda_1 > \mu_1 > 1$, then $n-2-\mu_1 > \mu_1$ and so the values for μ_1 are $2, 3, \ldots, \lceil \frac{n-2}{2} \rceil - 1$. Thus there are $q(q-1)$ choices for (α,β) and $\lceil \frac{n-2}{2} \rceil - 2 = \lceil \frac{n-6}{2} \rceil$ choices for μ_1 giving $q(q-1)\lceil \frac{n-6}{2} \rceil$ orbits. (b) If $\lambda_1 = \mu_1 > 1$, then n is even, $\lambda_1 = \mu_1 = \frac{n-2}{2}$, and there are $q(q-1)/2$ orbits as swapping α and β gives a matrix in the same orbit. (c) If $\lambda_1 > \mu_1 = 1$, then $\lambda_1 = n-3$ and there are $q(q-1)$ orbits. The number of orbits in cases (a) and (b) combined equals $q(q-1)\left(\frac{n-5}{2}\right)$ because if n is odd then $\lceil \frac{n-6}{2} \rceil = \frac{n-5}{2}$, while if n is even then $\lceil \frac{n-6}{2} \rceil + \frac{1}{2}$ also equals $\frac{n-5}{2}$. Thus the total number of matrices X in these three cases is:

$$\begin{aligned} &q(q-1)\left(\frac{n-5}{2}\right)\frac{q^{n^2-n-4}\omega(n,q)}{(1-q^{-1})^4} + q(q-1)\frac{q^{n^2-n-4}\omega(n,q)}{(1-q^{-1})^3(1-q^{-2})} \\ &= \frac{q^{n^2-n-2}\omega(n,q)}{(1-q^{-1})^3}\left[\frac{n-5}{2} + \frac{1}{1+q^{-1}}\right]. \end{aligned}$$

By Lemma 10, $\omega(n,q) > (1-q^{-1})^2$, and also $\frac{1}{1-q^{-1}} > 1+q^{-1}$. As $n \geqslant 5$ the above expression is greater than

$$q^{n^2-n-2}(1+q^{-1})\left[\frac{n-5}{2} + \frac{1}{1+q^{-1}}\right] > q^{n^2-n-2}\left(\frac{n-3}{2}\right). \tag{13}$$

The number of uncyclic matrices of type $(t-\alpha)^\lambda$ for some α is by Lemma 11 at least

$$q^{n^2-n-1}\left(1+q^{-1}-q^{-3}\right). \tag{14}$$

Adding the lower bound (14) to the lower bound (13) for the number of uncyclic matrices of type $(t-\alpha)^{(\lambda_1,1)}(t-\beta)^{(\mu_1,1)}$ gives the lower bound

$$\mathrm{unc}(n,q) > q^{n^2-n-1}\left(1 + \left(\frac{n-1}{2}\right)q^{-1} - q^{-3}\right)$$

of Theorem 9. □

5. An upper bound for $\mathrm{unc}(n,q)$ where $q>2$

It surprised the authors that mathematical induction, as employed in the proof of Theorem 14 below, could be used successfully to find an upper bound for $\mathrm{unc}(n,q)$ of the form postulated in Conjecture 3.

First we consider uncyclic matrices involving a unique irreducible f. Let $\mathrm{Irr}(r,q)$ denote the set of monic degree-r irreducible polynomials over $\mathbb{F}_q$. Recall that $N(r,q):=|\mathrm{Irr}(r,q)|$, and that $\omega(n,q):=\prod_{i=1}^{n}(1-q^{-i})=q^{-n^2}|\mathrm{GL}(n,q)|$.

Lemma 12. *Let $r(n,q)$ denote the cardinality of the set*

$$\{X\in \mathrm{M}(n,q)\mid X \text{ is uncyclic, and } c_X(t)=f^{n/d} \text{ for some } d|n, \text{ and some } f\in\mathrm{Irr}(d,q)\}$$

and set $c_1(n,q):=r(n,q)/q^{n^2-n-1}$. If $n\geqslant 2$, then $c_1(n,q)<c^*(q)$ where

$$c^*(q):=c_0(\infty,q)+q\frac{\omega(4,q)c_0(\infty,q^2)}{\omega(\infty,q^2)}\left(q\log(1-q^{-2})-\log(1-q^{-1})\right).$$

Moreover, $1+q^{-1}-q^{-3}<c^*(q)<1+\frac{3}{2}q^{-1}+\frac{2}{3}q^{-2}$ and $\lim_{q\to\infty}c^*(q)=1$.

Proof. It follows from the remarks preceding Lemma 11 that

$$r(n,q)=r(n,q,1)+\sum_{\substack{d|n\\1<d<n}} r(n,q,d) \tag{15}$$

because $r(n,q,n)=0$. Thus $r(n,q,1)\leqslant r(n,q)$ and so, by Lemma 11, $c_0(n,q)\leqslant c_1(n,q)$ with equality if and only if n is prime. It follows from Lemma 11 that $1+q^{-1}-q^{-3}=c_0(3,q)<c_0(\infty,q)<c^*(q)$. It remains to prove that $c_1(n,q)<c^*(q)$ for $n\geqslant 2$ and that $c^*(q)<1+\frac{3}{2}q^{-1}+\frac{2}{3}q^{-2}$. The first inequality is true when $n=2,3$ by Lemma 11 as

$$c_1(2,q)=c_0(2,q)=1<c_1(3,q)=c_0(3,q)=1+q^{-1}-q^{-3}<c_0(\infty,q)<c^*(q).$$

Assume henceforth that $n\geqslant 4$.

We digress to generalize the formula for $r(n,q,1)=c_0(n,q)q^{n^2-n-1}$ in Lemma 11 to $r(n,q,d)$. It follows from (2) and (3) that

$$r(n,q,1)=N(1,q)\sum_{\substack{\lambda\vdash n\\ \lambda\neq(n)}}\frac{|\mathrm{GL}(n,q)|}{c(\lambda,q)} \quad\text{and}\quad r(n,q,d)=N(d,q)\sum_{\substack{\lambda\vdash \frac{n}{d}\\ \lambda\neq(\frac{n}{d})}}\frac{|\mathrm{GL}(n,q)|}{c(\lambda,q^d)}$$

where the sums are over all partitions with more than one part. Note that the first sum counts the elements of the disjoint union $\bigcup_{\alpha\in\mathbb{F}_q} \mathrm{U}(n,q,\alpha)$. Relating these formulas gives

$$\begin{aligned} r(n,q,d) &= \frac{N(d,q)|\mathrm{GL}(n,q)|}{N(1,q^d)|\mathrm{GL}(\frac{n}{d},q^d)|}\, N(1,q^d) \sum_{\substack{\lambda\vdash\frac{n}{d}\\ \lambda\neq(\frac{n}{d})}} \frac{|\mathrm{GL}(\frac{n}{d},q^d)|}{c(\lambda,q^d)} \\ &= \frac{N(d,q)|\mathrm{GL}(n,q)|}{N(1,q^d)|\mathrm{GL}(\frac{n}{d},q^d)|} r(\frac{n}{d},q^d,1). \end{aligned}$$

By Lemma 11 we have $r(n,q,1) = c_0(n,q)q^{n^2-n-1}$, and so

$$\begin{aligned} r(n,q,d) &= \frac{q^{-d}N(d,q)q^{n^2}\omega(n,q)}{(q^d)^{(n/d)^2}\omega(\frac{n}{d},q^d)} c_0(\frac{n}{d},q^d)(q^d)^{(n/d)^2-(n/d)-1} \\ &= \frac{\omega(n,q)c_0(\frac{n}{d},q^d)}{\omega(\frac{n}{d},q^d)} \left[q^{-d}N(d,q)\right] q^{n^2-n-d} \end{aligned} \tag{16}$$

Since $n \geqslant 4$ and $1 < d < n$, each of d and n/d is at least 2, and so we have

$$\frac{\omega(n,q)c_0(\frac{n}{d},q^d)}{\omega(\frac{n}{d},q^d)} < \frac{\omega(4,q)c_0(\infty,q^2)}{\omega(\infty,q^2)} =: \gamma(q). \tag{17}$$

It follows from (15),(16) and (17) that

$$r(n,q) = r(n,q,1) + \sum_{\substack{d|n\\ 1<d<n}} r(n,q,d) \leqslant c_0(\infty,q)q^{n^2-n-1} + \gamma(q)\left(\sum_{\substack{d|n\\ 1<d<n}} q^{-d}N(d,q)q^{n^2-n-d}\right).$$

This proves that $r(n,q) \leqslant K(n,q)q^{n^2-n-1}$ where

$$K(n,q) := c_0(\infty,q) + q\gamma(q) \sum_{\substack{d|n\\ 1<d<n}} q^{-2d}N(d,q).$$

Thus $c_1(n,q) \leqslant K(n,q)$, and our goal now is to prove that $K(n,q) < c^*(q)$ for $n \geqslant 4$.

The bound $N(d,q) \leqslant (q^d - q)/d$, which holds for $d \geqslant 2$, gives

$$\sum_{\substack{d|n\\ 1<d<n}} q^{-2d}N(d,q) \leqslant \sum_{d\geqslant 2} \frac{q^{-d}}{d} - q\sum_{d\geqslant 2}\frac{q^{-2d}}{d} = \sum_{d\geqslant 1}\frac{q^{-d}}{d} - q\sum_{d\geqslant 1}\frac{q^{-2d}}{d}. \tag{18}$$

The series $\sum_{d\geqslant 1}\frac{x^d}{d}$ converges absolutely for $|x| < 1$ to $-\log(1-x)$. Thus

$$K(n,q) < c_0(\infty,q) + q\gamma(q)\left(q\log(1-q^{-2}) - \log(1-q^{-1})\right) = c^*(q),$$

so $c_1(n,q) \leqslant K(n,q) < c^*(q)$ for $n \geqslant 4$. Finally we must show that $c^*(q) < 1 + \frac{3}{2}q^{-1} + \frac{2}{3}q^{-2}$.

We begin by showing $q\log(1-q^{-2})-\log(1-q^{-1})<q^{-2}/2$ for $q\geqslant 2$. This is true when $q=2$ because $2\log(3/4)-\log(1/2)<0.125$. Suppose now that $q\geqslant 3$. If $0\leqslant x<1$, then elementary calculus gives

$$x+\frac{x^2}{2}\leqslant -\log(1-x)\leqslant x+\frac{x^2}{2}+\sum_{d\geqslant 3}\frac{x^d}{3}=x+\frac{x^2}{2}+\frac{x^3}{3(1-x)}.$$

If $x=q^{-2}$, then $q^{-2}+\frac{q^{-4}}{2}\leqslant -\log(1-q^{-2})$ and $q\log(1-q^{-2})\leqslant -q^{-1}-\frac{q^{-3}}{2}$. If $x=q^{-1}$, then $-\log(1-q^{-1})<q^{-1}+\frac{q^{-2}}{2}+\frac{q^{-3}}{3(1-q^{-1})}\leqslant q^{-1}+\frac{q^{-2}}{2}+\frac{q^{-3}}{2}$ for $q\geqslant 3$. Adding shows

$$q\log(1-q^{-2})-\log(1-q^{-1})<\frac{q^{-2}}{2}\qquad\text{for } q\geqslant 3. \tag{19}$$

By Lemma 13 below, $c^*(2)<1.83<\frac{23}{12}$, and hence the bound $c^*(q)<1+\frac{3}{2}q^{-1}+\frac{2}{3}q^{-2}$ holds when $q=2$. Assume henceforth that $q\geqslant 3$. Lemma 11 gives $c_0(\infty,q)\leqslant 1+q^{-1}+q^{-2}$, and hence $c_0(\infty,q^2)\leqslant 1+q^{-2}+q^{-4}$. Lemma 10 implies $\omega(\infty,q^2)>1-q^{-2}-q^{-4}$, and Lemma 7 may be used to show that $\omega(\infty,q^2)^{-1}<1+q^{-2}+3q^{-4}$ for $q\geqslant 3$. The inequalities $\omega(4,q)<(1-q^{-1})(1-q^{-2})$ and (19) give:

$$\begin{aligned}c^*(q)&<(1+q^{-1}+q^{-2})+(1-q^{-1})(1-q^{-2})(1+q^{-2}+q^{-4})(1+q^{-2}+3q^{-4})\frac{q^{-1}}{2}\\&=1+\frac{1}{2}\left(3q^{-1}+q^{-2}+q^{-3}-q^{-4}+3q^{-5}-3q^{-6}-q^{-7}+q^{-8}-q^{-9}+q^{-10}-3q^{-11}+3q^{-12}\right)\\&\leqslant 1+\frac{3}{2}q^{-1}+\frac{2}{3}q^{-2},\end{aligned}$$

where the final inequality follows from Lemma 7 with $q_0=3$. As q approaches infinity, the established lower and upper bounds for $c^*(q)$ both approach 1. Thus $\lim_{q\to\infty}c^*(q)=1$ as claimed. This completes the proof. □

The proof of our main theorem requires sharper bounds for $c^*(2)$ and $c^*(3)$ than those provided by Lemma 12.

Lemma 13. *For $m\geqslant 2$, $q\geqslant 2$, we have*

$$\omega(\infty,q)>\omega(m-1,q)\left(1-\frac{q^{-m}}{1-q^{-1}}\right) \tag{20}$$

and this bound may be used to show that $c^*(2)<1.83$ and $c^*(3)<1.56$.

Proof. The bound $\prod_{i=m}^{\infty}(1-q^{-i})>1-\sum_{i=m}^{\infty}q^{-i}$ gives rise to the lower bound (20) for $\omega(\infty,q)$. This, in turn, gives an upper bound for $c_0(\infty,q)$ (see Lemma 11 for a definition). Setting $m=6$ and $q=2,4$ in (20) gives

$$\omega(\infty,2)>0.28869,\quad \omega(\infty,4)>0.688,\quad c_0(\infty,2)<1.691,\quad\text{and}\quad c_0(\infty,4)<1.312.$$

Similarly, setting $m=4$ and $q=3,9$ in (20) gives

$$\omega(\infty,3)>0.560,\quad \omega(\infty,9)>0.876,\quad c_0(\infty,3)<1.439,\quad\text{and}\quad c_0(\infty,9)<1.124.$$

These inequalities give $c^*(2) < 1.83$ and $c^*(3) < 1.56$. □

Theorem 14. *If $n \geqslant 1$, then $\operatorname{unc}(n,3) < (1.56)3^{n^2-n-1}\,(1.59)^n$ and*

$$\operatorname{unc}(n,q) < c^*(q)q^{n^2-n-1}\left(1+q^{-1}+2q^{-2}\right)^n \tag{21}$$

for $q \geqslant 4$, where $c^*(q)$ is defined in Lemma 12*, and satisfies $1 < c^*(q) < 1.56$ for $q \geqslant 3$.*

Proof. Our proof has two parts. First, we use induction on n and a geometric argument to prove $\operatorname{unc}(n,q) \leqslant c^*(q)q^{n^2-n-1}\rho(q)^n$ for $n \geqslant 1$ and $q \geqslant 2$, where

$$\rho(q) := \frac{1+\sqrt{1+\frac{4c^*(q)}{q\omega(\infty,q)}}}{2}. \tag{22}$$

Second, we prove that $\rho(3) < 1.59$, and $\rho(q) < 1+q^{-1}+2q^{-2}$ for $q \geqslant 4$.

It follows from the definition (22) that $\rho(q) > 1$ for all $q \geqslant 2$. A simple calculation shows that $\operatorname{unc}(n,q) \leqslant c^*(q)q^{n^2-n-1}\rho(q)^n$ is true for $n = 1, 2$ and all q. Consider the proof for $n = 3$. By Table 3, $\operatorname{unc}(3,q) = q^5+q^4-q^2$ and so the inequality to be proved is:

$$q^5\left(1+q^{-1}-q^{-3}\right) \leqslant c^*(q)q^5\rho(q)^3.$$

Now by Lemma 12, $1+q^{-1}-q^{-3} < c^*(q)$, and as $\rho(q) > 1$ the inequality above holds for all q. Assume henceforth that $n \geqslant 4$.

By definition, there are precisely $r(n,q)$ uncyclic matrices $X \in \mathrm{M}(n,q)$ for which $c_X(t)$ is a power of some irreducible polynomial. We shall now over-estimate the number of uncyclic X for which $c_X(t)$ is *not* a power of a single irreducible.

We impose an arbitrary total ordering on the (finite number of) irreducible polynomials over $\mathbb{F}_q$ of degree at most n. For each uncyclic matrix X such that $c_X(t)$ is not a power of an irreducible, there exists at least one irreducible polynomial f such that, if $f^{\nu(f)}$ is the highest power of f dividing $c_X(t)$, then $0 < d(f)\nu(f) \leqslant n/2$. We choose the first irreducible f in the total ordering with this property. Write $V = U \oplus W$, where $U = \ker f(X)^{\nu(f)}$ is the f-primary component and $W = \operatorname{im} f(X)^{\nu(f)}$ is an X-invariant complement. The restrictions X_U and X_W of X to U and W are both uncyclic. Moreover, X determines a unique 4-tuple (U, W, X_U, X_W). Counting the number of possible 4-tuples will give an upper bound for the number of X.

Set $k := \dim(U)$. Then $k = d(f)\nu(f) \leqslant n/2$, and $k \geqslant 2$ as X_U is uncyclic. The number of decompositions $V = U \oplus W$ with $\dim(U) = k$ is

$$\frac{|\mathrm{GL}(n,q)|}{|\mathrm{GL}(k,q)|\,|\mathrm{GL}(n-k,q)|}.$$

The number of choices for X_U is precisely $r(k,q)$, and the number of choices for X_W is at most $\operatorname{unc}(n-k,q)$. (At this point the reader may be concerned that we are not using the fact that the characteristic polynomial of X_W is coprime to f. It is remarkable

that this otherwise very delicate counting problem is essentially insensitive to such an over-estimation.) Thus

$$\mathrm{unc}(n,q) \leqslant r(n,q) + \sum_{k=2}^{\lfloor \frac{n}{2} \rfloor} \frac{|\mathrm{GL}(n,q)|}{|\mathrm{GL}(k,q)|\,|\mathrm{GL}(n-k,q)|}\, r(k,q)\,\mathrm{unc}(n-k,q).$$

We shall abbreviate $\rho(q)$, $c^*(q)$ and $\omega(\infty,q)$ by ρ, c^* and ω, respectively. As $n-k<n$, it follows by induction that

$$\mathrm{unc}(n-k,q) \leqslant c^* q^{(n-k)^2-(n-k)-1}\rho^{n-k}.$$

Moreover, Lemma 12 gives $r(k,q) = c_1(k,q)q^{k^2-k-1} \leqslant c^*q^{k^2-k-1}$ for all k and, since $\frac{\omega(n,q)}{\omega(n-k,q)} = \prod_{i=n-k+1}^{n}(1-q^{-i}) < 1$, we have

$$\frac{|\mathrm{GL}(n,q)|}{|\mathrm{GL}(k,q)|\,|\mathrm{GL}(n-k,q)|} = \frac{\omega(n,q)q^{n^2-k^2-(n-k)^2}}{\omega(k,q)\omega(n-k,q)} < \frac{q^{n^2-k^2-(n-k)^2}}{\omega(k,q)}.$$

Thus

$$\mathrm{unc}(n,q) \leqslant c^*q^{n^2-n-1} + \sum_{k=2}^{\lfloor \frac{n}{2} \rfloor} \frac{q^{n^2-k^2-(n-k)^2}}{\omega(k,q)} c_1(k,q)q^{k^2-k-1}\, c^*q^{(n-k)^2-(n-k)-1}\rho^{n-k}.$$

The exponent of q in the terms of the summation is independent of k as

$$n^2-k^2-(n-k)^2+k^2-k-1+(n-k)^2-(n-k)-1 = n^2-n-2.$$

Therefore

$$\mathrm{unc}(n,q) \leqslant c^*q^{n^2-n-1}\left(1+\sum_{k=2}^{\lfloor \frac{n}{2} \rfloor} \frac{c_1(k,q)}{q\omega(k,q)}\rho^{n-k}\right). \tag{23}$$

To complete the induction we must show that the above bracketed expression is at most ρ^n. Towards this end, note that $\frac{c_1(k,q)}{\omega(k,q)} < \frac{c^*}{\omega}$ by Lemma 12. Since $\lfloor \frac{n}{2} \rfloor + \lceil \frac{n}{2} \rceil = n$, $n \geqslant 4$ and $\rho > 1$, we have

$$\sum_{k=2}^{\lfloor \frac{n}{2} \rfloor} \rho^{n-k} = \rho^{\lceil \frac{n}{2} \rceil} + \rho^{\lceil \frac{n}{2} \rceil+1} + \cdots + \rho^{n-2} = \frac{\rho^{\lceil \frac{n}{2} \rceil}(\rho^{\lfloor \frac{n}{2} \rfloor-1}-1)}{\rho-1} \leqslant \frac{\rho^{n-1}-\rho^2}{\rho-1}.$$

It follows from the definition (22) of ρ, by rationalizing the denominator, that

$$\frac{1}{\rho-1} = \frac{2}{-1+\sqrt{1+\frac{4c^*}{q\omega}}} = \frac{q\omega}{2c^*}\left[1+\sqrt{1+\frac{4c^*}{q\omega}}\right] = \frac{q\omega\rho}{c^*}.$$

The previous three displayed equations now give

$$\mathrm{unc}(n,q) \leqslant c^*q^{n^2-n-1}\left(1+\frac{c^*}{q\omega}\frac{q\omega\rho}{c^*}(\rho^{n-1}-\rho^2)\right) = c^*q^{n^2-n-1}\left(1+\rho^n-\rho^3\right).$$

Since $\rho > 1$, it follows that $1 - \rho^3 < 0$. Thus $\mathrm{unc}(n,q) < c^* q^{n^2-n-1}\rho^n$ and we have completed the inductive proof.

To complete the proof, we must estimate $\rho(q)$. By Lemma 13, $c^*(3) < 1.56$ and $\omega(\infty, 3) < 0.56$. Thus $\rho(3) < 1.59$, and the inequality for $\mathrm{unc}(n,3)$ follows. Assume now that $q \geqslant 4$. We will show that

$$\rho(q) = \frac{1 + \sqrt{1 + \frac{4c^*}{q\omega}}}{2} < 1 + q^{-1} + 2q^{-2}. \tag{24}$$

Multiplying (24) by 2, subtracting 1, and squaring gives that (24) is equivalent to

$$1 + \frac{4c^*}{q\omega} < (1 + 2q^{-1} + 4q^{-2})^2 = 1 + 4q^{-1} + 12q^{-2} + 16q^{-3} + 16q^{-4}. \tag{25}$$

Subtracting 1 from (25) and multiplying by the positive quantity $\frac{q\omega}{4}$ gives the equivalent inequality

$$c^* < \omega(1 + 3q^{-1} + 4q^{-2} + 4q^{-3}). \tag{26}$$

By virtue of the inequalities $c^* < 1 + \frac{3}{2}q^{-1} + \frac{2}{3}q^{-2}$ from Lemma 12, and $1 - q^{-1} - q^{-2} < \omega$ from Lemma 10, the inequality (26), and hence also the required equivalent inequality (24), will follow from a proof of the following stronger inequality:

$$1 + \frac{3}{2}q^{-1} + \frac{2}{3}q^{-2} < (1 - q^{-1} - q^{-2})(1 + 3q^{-1} + 4q^{-2} + 4q^{-3}). \tag{27}$$

Expanding and rearranging (27) gives

$$0 < \frac{q^{-1}}{2} - \frac{2q^{-2}}{3} - 3q^{-3} - 8q^{-4} - 4q^{-5}. \tag{28}$$

This inequality is true for $q \geqslant 4$ by Lemma 7 with $q_0 = 4$. Thus (24) holds for $q \geqslant 4$. This completes the proof. □

Corollary 15. *If $n \geqslant 1$ and $q \geqslant 3$, then the probability p that a uniformly distributed random $n \times n$ matrix over $\mathbb{F}_q$ is f-cyclic satisfies*

$$1 - k(q)q^{-1}\left(q^{-1} + q^{-2} + \frac{12}{5}q^{-3}\right)^n < p \leqslant 1 \tag{29}$$

where $k(q) = 1 + \frac{3}{2}q^{-1} + \frac{2}{3}q^{-2}$.

Proof. Note that $p = 1 - \mathrm{unc}(n,q)q^{-n^2} \leqslant 1$. Theorem 14 with $q = 3$ gives

$$\frac{\mathrm{unc}(n,3)}{3^{n^2}} < \frac{1.56}{3}\left(\frac{1.59}{3}\right)^n < \frac{k(3)}{3}\left(3^{-1} + 3^{-2} + \frac{12 \cdot 3^{-3}}{5}\right)^n.$$

Thus the lower bound for p in (29) holds for $q = 3$. Assume now that $q \geqslant 4$. Since $c^*(q) < k(q)$ by Lemma 12, it follows from Theorem 14 that

$$\frac{\mathrm{unc}(n,q)}{q^{n^2}} < k(q)q^{-1}\left(q^{-1} + q^{-2} + 2q^{-3}\right)^n < k(q)q^{-1}\left(q^{-1} + q^{-2} + \frac{12}{5}q^{-3}\right)^n.$$

This establishes the lower bound for p in (29) for $q \geqslant 4$, and completes the proof. □

6. An upper bound for $\mathrm{unc}(n, 2)$

Theorem 14 shows that $\mathrm{unc}(n, q)/q^{n^2} = \mathrm{O}(R(q)^n)$, where $R(q) = \rho(q)/q$ with $\rho(q)$ as defined in (22). For this value of $\rho(q)$, the proof of Theorem 14 yields an upper bound for $\rho(q)$, and hence also for $R(q) = \rho(q)/q$, as listed in Table 4, for various values of q. (The values of these bounds have been rounded up to the nearest 10^{-2}.) We note that the inductive part of the proof of Theorem 14 is valid for $q = 2$, but it gives an upper bound for $R(2)$ greater than 1, or equivalently for $\rho(2)$ greater than 2. Stronger arguments are needed to show that $\mathrm{unc}(n, 2)/2^{n^2} = \mathrm{O}(R(2)^n)$ with $R(2) < 1$. If Conjecture 3 were true, then this would hold with $R(2) \leqslant \frac{1}{2} + \frac{1}{2\times 2^2} = 0.625$ (and hence with $\rho(2) = 2R(2) \leqslant 1.25$). In this section we modify the proof of Theorem 14 to obtain a value of $R(2)$ less than 0.983, or $\rho(2)$ less than 1.966, which is still substantially larger than the bound conjectured to hold in Conjecture 3. Theorem 16 below implies Theorem 1.

q	2	3	4	5	7	8	9	11	13	16	17	19	23
$R(q) \leqslant$	1.18	0.53	0.35	0.26	0.15	0.17	0.13	0.11	0.09	0.07	0.07	0.06	0.05
$\rho(q) \leqslant$	2.35	1.59	1.38	1.28	1.18	1.16	1.14	1.11	1.09	1.07	1.07	1.06	1.05

Table 4. Upper bounds for $R(q) = \rho(q)/q$ and $\rho(q)$ obtained in Theorem 14.

Theorem 16. *If $n \geqslant 1$, then* $2^{n^2-n-1}(\frac{n}{4} + \frac{5}{8}) < \mathrm{unc}(n, 2) < (1.83)2^{n^2-n-1}(1.966)^n$.

Proof. The lower bound follows from Theorem 9. The upper bound is proved by adapting the inductive proof of Theorem 14. By Proposition 8 we know that $\mathrm{unc}(n, 2)$ is at most $2^{n^2-n-1}(1.25)^n$ for $n \leqslant 9$ (indeed even for $n \leqslant 37$), so the weaker bound $\mathrm{unc}(n, 2) < (1.83)2^{n^2-n-1}(1.966)^n$ certainly holds for $n \leqslant 9$. Assume henceforth that $n \geqslant 10$. Lemma 13 shows that $c^*(2)$ as defined in Lemma 12 satisfies $c^*(2) < 1.83$. Set $\rho := 1.966$. The first part of the proof of Theorem 14 is valid for $q = 2$, and in particular, the inequality (23) holds for $q = 2$. To complete the inductive step in the proof it is sufficient to prove, for $n \geqslant 10$, that

$$1 + \sum_{k=2}^{\lfloor \frac{n}{2} \rfloor} \frac{c_1(k, 2)}{2\omega(k, 2)} \rho^{n-k} \leqslant \rho^n \qquad \text{or equivalently,} \qquad \rho^{-n} + \sum_{k=2}^{\lfloor \frac{n}{2} \rfloor} \frac{c_1(k, 2)}{2\omega(k, 2)} \rho^{-k} \leqslant 1$$

with $c_1(k, 2)$ as defined in Lemma 12. Since $\rho^{-n} \leqslant \rho^{-10}$ it is sufficient to prove that

$$\rho^{-10} + \sum_{k=2}^{\infty} \frac{c_1(k, 2)}{2\omega(k, 2)} \rho^{-k} \leqslant 1. \tag{30}$$

For $k \geqslant 6$ we use the bounds from Lemmas 12 and 13 to obtain

$$\frac{c_1(k,2)}{2\omega(k,2)} < \frac{c^*(2)}{2\omega(\infty,2)} < \frac{1.83}{2 \times 0.28869} < 3.17. \tag{31}$$

For $k \leqslant 5$ we use the exact values of $\frac{c_1(k,2)}{2\omega(k,2)}$. Recall from the definitions of $c_0(k,q)$ and $c_1(k,q)$ in Lemmas 11 and 12 that $c_1(k,q)$ equals $c_0(k,q)$ when k is prime. Hence $c_1(k,2)$ equals $1, \frac{11}{8}, \frac{6619}{4098}$ for $k = 2,3,5$. To compute $c_1(4,2)$, we use the proof of Lemma 12 to show $r(4,2) = r(4,2,1) + r(4,2,2) = 3152 + 112 = 3264$. Thus $c_1(4,2) = r(4,2)/2^{11} = \frac{51}{32}$.

k	2	3	4	5
$\frac{c_1(k,2)}{2\omega(k,2)}$	$\frac{4}{3}$	$\frac{44}{21}$	$\frac{272}{105}$	$\frac{26476}{9765}$

TABLE 5. Values of $\frac{c_1(k,2)}{2\omega(k,2)}$ for $2 \leqslant k \leqslant 5$.

Using (31) and Table 5, the infinite sum in (30) is less than

$$\sum_{k=2}^{5} \frac{c_1(k,2)}{2\omega(k,2)}\rho^{-k} + \frac{c^*(2)}{2\omega(\infty,2)} \sum_{k=6}^{\infty} \rho^{-k} < \frac{4}{3}\rho^{-2} + \frac{44}{21}\rho^{-3} + \frac{272}{105}\rho^{-4} + \frac{26476}{9765}\rho^{-5} + \frac{3.17\rho^{-6}}{1-\rho^{-1}}.$$

Evaluating the expression

$$\rho^{-10} + \frac{4}{3}\rho^{-2} + \frac{44}{21}\rho^{-3} + \frac{272}{105}\rho^{-4} + \frac{26476}{9765}\rho^{-5} + \frac{3.17\rho^{-6}}{1-\rho^{-1}}$$

at $\rho = 1.966$ gives the number $0.9992\cdots < 1$. This completes the inductive proof. □

7. Finding a witness to X being f-cyclic

In this section h always denotes a *monic* irreducible polynomial. Henceforth we shall consistently omit the adjective "monic". The h-primary component $V(h)$ of an $\mathbb{F}_q[X]$-module V can be generalized to $V(g)$ where g is a (possibly reducible) divisor of $c_X(t)$: set $V(g) := \bigoplus_{h|g} V(h)$ where the sum is over irreducible divisors h of g.

The Holt-Rees MEAT-AXE algorithm [16, Section 2] initially finds a random matrix X, and then begins by executing the following steps:

(1) find an irreducible factor g of the characteristic polynomial $c_X(t)$,
(2) evaluate $g(t)$ at X to compute $Y = g(X)$, and
(3) find a non-zero vector $u \in \ker(Y)$.

The matrix X can be used to prove irreducibility if it is f-cyclic relative to g, that is, if (and only if) the degree of g equals $\dim(\ker(Y))$. Step (2) has cost $\mathrm{O}(\mathrm{Mat}(n)n)$ field

operations[1] (that is, additions, subtractions, multiplications, and inversions in $\mathbb{F}_q$), where $\mathrm{Mat}(n)$ is an upper bound for the number of field operations required to multiply two matrices in $\mathrm{M}(n,q)$. The purpose of this section is to present a one-sided Monte Carlo algorithm called Is*f*Cyclic that requires (only) $\mathrm{O}(\mathrm{Mat}(n)\log n)$ field operations, and in particular obviates the necessity of applying the rather expensive Step (2).

Given an f-cyclic matrix $X \in \mathrm{M}(n,q)$, and a positive real number $\varepsilon < 1$, this algorithm returns True with probability at least $1-\varepsilon$. Moreover in this case it constructs a divisor g of $c_X(t)$ and a non-zero vector u such that $\gcd(g, c_X/g) = 1$ and $V(g) = u\mathbb{F}_q[X]$. This shows that X is f-cyclic relative to every irreducible divisor of g. If Is*f*Cyclic fails to construct g, u with these properties then it returns False, that is to say, Is*f*Cyclic incorrectly reports 'X is not f-cyclic'. However, the probability of this happening is at most ε. On the other hand, if X is not f-cyclic, then Is*f*Cyclic correctly returns False. In summary, an output True is always correct, while an output False is incorrect with probability at most ε. These assertions are proved in Theorem 18.

If it were desirable that the polynomial g returned by the algorithm Is*f*Cyclic be irreducible, then Is*f*Cyclic could be modified to incorporate a randomised polynomial factorisation algorithm.

7.1. **Witnesses and orders.** Given a matrix $X \in \mathrm{M}(n,q)$ and a non-constant divisor g of $c_X(t) = \prod_f f^{\nu(f)}$, a vector $v \in V := \mathbb{F}_q^{1\times n}$ is called a *g-witness* for X if the cyclic submodule $v\mathbb{F}_q[X]$ contains the h-primary component $V(h)$ of V for all irreducible divisors h of g. The following are equivalent: (1) v is a g-witness for X, (2) $V(g) \subseteq v\mathbb{F}_q[X]$, and (3) $\prod_h h^{\nu(h)}$ divides the order polynomial $\mathrm{ord}_X(v)$, where the product is over all irreducible divisors h of g. (Recall that $a(t) = \mathrm{ord}_X(v)$ is the smallest degree monic polynomial over $\mathbb{F}_q$ satisfying $va(X) = 0$.) As submodules of cyclic modules are cyclic, *X has a g-witness v if and only if X is f-cyclic relative to every irreducible divisor h of g.*

It turns out that a matrix X, which is f-cyclic relative to every irreducible divisor of g, has many g-witnesses, and failure to find a g-witness (for any such g) provides "probabilistic evidence" that X is uncyclic (as is shown below).

Recall the following notation from Section 2

$$\mathrm{type}(X) = \prod_h h^{\lambda(h)}, \quad c_X(t) = \prod_h h^{|\lambda(h)|}, \quad m_X(t) = \prod_h h^{\lambda(h)_1}, \quad V(h) = \ker h(X)^{\lambda(h)_1},$$

and set $V(h)_k := \ker h(X)^{\lambda(h)_1 - k}$ for $0 \leqslant k \leqslant \lambda(h)_1$. The subspaces $V(h)_k$ define a chain

$$V(h) = V(h)_0 > V(h)_1 > \cdots > V(h)_{\lambda(h)_1} = 0 \tag{32}$$

of $\mathbb{F}_q[X]$-submodules.

[1] A lower complexity can be achieved by conjugating X into Frobenius normal form, evaluating $g(t)$ at the matrix obtained, and conjugating back. For the complexity of this approach see: C. Pernet and A. Storjohann, Frobenius form in expected matrix multiplication time over sufficiently large fields, preprint.

We introduce the notion of the *h-order* of a vector or polynomial, see [10, 7.17]. Fix an irreducible polynomial $h(t)$, and let I be the ideal $h(t)\mathbb{F}_q[t]$ of $\mathbb{F}_q[t]$. A non-zero vector v in an $\mathbb{F}_q[t]$-module M is said to have h-order k, written $o_h(v) = k$, if k is the largest integer such that $v \in MI^k$. By convention we set $o_h(0) := \infty$. In our applications, the module M will be either V, the h-primary component $V(h)$, or the ring $\mathbb{F}_q[t]$. We denote elements of V by u, v, and elements of $\mathbb{F}_q[t]$ by a, d, e, g. In the case when $M = \mathbb{F}_q[t]$, we have $\cap_{k\geqslant 0} MI^k = 0$, and o_h is an exponential valuation satisfying: (i) $o_h(a) = \infty$ if and only if $a = 0$, (ii) $o_h(ab) = o_h(a) + o_h(b)$, (iii) $o_h(a + b) \geqslant \min(o_h(a), o_h(b))$, and (iv) $o_h(\gcd(a, b)) = \min(o_h(a), o_h(b))$. When $M = V(h)$ properties (i) and (iii) hold.

Suppose that $v \in V(h)$. Then $o_h(v) = k$ holds if $v \in V(h)_k$ and k is maximal. If $v \neq 0$, then $o_h(\mathrm{ord}_X(v)) \leqslant \nu(h) - o_h(v)$, and $\dim_{\mathbb{F}_q}(V(h)/V(h)_k) \geqslant k\deg(h)$ for all $k \leqslant \lambda(h)_1$. These inequalities become equalities when X is f-cyclic relative to h. In the case that X is f-cyclic relative to h, then $V(h)$ is uniserial, and a uniformly distributed random vector $v \in V(h)$ has $o_h(v) = k$ with probability

$$\frac{|V(h)_k| - |V(h)_{k+1}|}{|V(h)|} = q^{-k\deg(h)} - q^{-(k+1)\deg(h)}.$$

Each vector $v \in V$ has a unique decomposition $v = \sum_h v_h$ where each h is irreducible and v_h belongs to the h-primary component $V(h)$ of V. Thus, for a non-constant divisor g of $c_X(t)$, v is a g-witness if and only if $v_h \notin V(h)_1$ holds for each irreducible divisor h of g, or equivalently, $o_h(v) = 0$ for each irreducible divisor h of g. This happens with probability $\prod_{h|g}(1 - q^{-\deg(h)})$, where the product is over all (monic) irreducible divisors h of g.

7.2. **Is*f*Witness.** The algorithm Is*f*Cyclic has input (X, ε), and makes repeated calls to a deterministic subprogram Is*f*Witness with input $(v, X, c_X(t))$, where v is a uniformly distributed random vector in $V = \mathbb{F}_q^{1\times n}$. Because $c_X(t)$ should be calculated once, and not each time the subprogram Is*f*Witness is invoked, it is listed as an input parameter for Is*f*Witness. The algorithm Is*f*Witness outputs True if v is a g-witness for X for some non-constant divisor g of $c_X(t)$, or False if v is not an g-witness for any non-constant divisor g of $c_X(t)$. As the Meat-axe requires a useful certificate of f-cyclicity, in the former case, Is*f*Witness outputs a triple $(\text{True}, u, a(t))$ where $u \neq 0$, $\mathrm{ord}_X(u) = a(t)$, $\gcd(a(t), c_X(t)/a(t)) = 1$, and u is an $a(t)$-witness.

The subprogram Is*f*Witness introduces a vector u and polynomials a, d, g that are modified in the course of the algorithm. However, each time line 5 is executed, the relations $u = vg(X)$, $a = \mathrm{ord}_X(u)$, and $d = \gcd(a, c_X(t)/a)$ always hold, see Theorem 17(a). It is useful to note that if d divides $a = \mathrm{ord}_X(u)$, then $\mathrm{ord}_X(ud(X)) = a/d$.

Algorithm. Is*f*Witness
Input. a non-zero vector $v \in V$; $X \in \mathrm{M}(n, q)$; the characteristic polynomial $c_X(t)$
Output. $(\text{True}, u, a(t))$, or False

1. $u := v$; $g(t) := 1$; # $u = vg(X)$ always holds

2. $a(t) := \mathrm{ord}_X(u)$; # compute the order polynomial of u under X
3. $d(t) := \gcd(a(t), c_X(t)/a(t))$; # d is always $\gcd(a, c_X(t)/a)$
4. $i := 1$;
5. WHILE $i \leqslant \lfloor \log_2 n \rfloor + 2$ do
6. if $d = 1$ then return $(\textsc{True}, u, a(t))$; fi;
7. if $d = a$ then return False; fi; # henceforth $d \neq 1, a$ and d divides a
8. $g := g * d$; $u := vg(X)$; # $u := ud(X)$ is less efficient
9. $a := a/d$; # $a = \mathrm{ord}_X(u) = \mathrm{ord}_X(v)/g$ always hold
10. $e := \gcd(a, d)$; $d := e * \gcd(a/e, e)$; # $d = \gcd(a, c_X(t)/a)$ always holds
11. $i := i+1$; # $i =$ number of times line 5 is executed

Theorem 17. *Parts* (a)–(e) *below prove the correctness of the algorithm* IsfWitness. *Let* $c_X(t) = \prod_h h^{\nu(h)}$, *where the product is over all (monic) irreducible divisors of* $c_X(t)$. *Suppose that line* 5 *is executed* s *times, and the values of* $u, a(t), d(t)$ *and* $g(t)$ *at the* i*th iteration of line 5 are* $u_i, a_i(t), d_i(t)$ *and* $g_i(t)$, *respectively. Also set* $b_i := c_X(t)/a_i(t)$.

(a) *Then* $u_i \neq 0$, $a_i = \mathrm{ord}_X(u_i) \neq 1$, $u_i = vg_i(X)$, *and* $d_i = \gcd(a_i, b_i)$ *for* $1 \leqslant i \leqslant s$.
(b) *Set* $k(h) := \nu(h) - o_h(\mathrm{ord}_X(v))$ *for each irreducible divisor* h *of* $c_X(t)$, *and set* $r(h) := \left\lfloor \log_2 \frac{\nu(h)}{k(h)} \right\rfloor + 1$ *when* $k(h) > 0$. *Then either*
  (i) $k(h) = 0$, *and for* $i \geqslant 1$, $o_h(a_i) = \nu(h)$ *and* $o_h(d_i) = 0$; *or*
  (ii) $k(h) > 0$, *and* $o_h(a_i) = o_h(d_i) = 0$ *for* $i \geqslant r(h) + 1$.
(c) *Set* $r := 0$ *if* $k(h) = 0$ *for all irreducibles* h, *and set* $r := \max\{r(h) \mid k(h) > 0\}$ *otherwise. Then* $s \leqslant r+1 \leqslant \lfloor \log_2 n \rfloor + 2$. *Also* IsfWitness *returns* $(\textsc{True}, u_s, a_s)$ *at line* 6, *or* False *at line* 7. *In either case,* $\max\{1, r-1\} \leqslant s \leqslant r+1$ *holds.*
(d) IsfWitness *returns* True *if and only if* v *is an* a*-witness for* X *for some non-constant divisor* a *of* $c_X(t)$.
(e) IsfWitness *returns* False *if and only if* $o_h(\mathrm{ord}_X(v)) < \nu(h)$ *for each irreducible polynomial* h *such that* X *is* f*-cyclic relative to* h. *In particular,* IsfWitness *returns* False *if* X *is uncyclic.*
(f) IsfWitness *requires* $\mathrm{O}(\mathrm{Mat}(n) \log n)$ *field operations.*

Proof. (a) We use induction on i. Part (a) holds for $i = 1$ by the definitions of u, a, d, g in lines 1–3 of IsfWitness. Suppose inductively that the claimed relations hold for $1 \leqslant i < s$. As $i < s$, IsfWitness does not terminate at lines 6 or 7 on the ith iteration, and it follows that $d_i \neq 1, a_i$. The new values of these variables assigned during the ith iteration of lines 8–10 are $g_{i+1} = g_i * d_i, u_{i+1} = vg_{i+1}(X), a_{i+1} = a_i/d_i$, and $d_{i+1} = e * \gcd(a_{i+1}/e, e)$ where $e = \gcd(a_{i+1}, d_i)$. Since $\mathrm{ord}_X(u_{i+1}) = \mathrm{ord}_X(vg_{i+1}) = \mathrm{ord}_X(u_i d_i) = a_i/d_i \neq 1$, it follows that $u_{i+1} \neq 0$ and $a_{i+1} = \mathrm{ord}_X(u_{i+1})$. By definition $b_i = c_X/a_i$, and hence $b_{i+1} = c_X/a_{i+1} = b_i * a_i/a_{i+1} = b_i * d_i$. Finally, we must prove that $d_{i+1} = \gcd(a_{i+1}, b_{i+1})$. Now $d_i = \gcd(a_i, b_i)$ implies that $\gcd(a_i/d_i, b_i/d_i) = 1$, that is, $\gcd(a_{i+1}, b_i/d_i) = 1$. Similarly $e = \gcd(a_{i+1}, d_i)$ implies that $\gcd(a_{i+1}/e, d_i/e) = 1$. To complete the inductive proof of part (a) we show that $\gcd(a_{i+1}, b_{i+1})$ is equal to

$e * \gcd(a_{i+1}/e, e)$, which is d_{i+1}:

$$\begin{aligned}
\gcd(a_{i+1}, b_{i+1}) &= \gcd\left(a_{i+1}, b_i * d_i\right) && \text{(since } b_{i+1} = b_i * d_i\text{)} \\
&= \gcd(a_{i+1}, \frac{b_i}{d_i} * d_i^2) \\
&= \gcd(a_{i+1}, d_i^2) && \text{(since } \gcd(a_{i+1}, \frac{b_i}{d_i}) = 1\text{)} \\
&= e * \gcd(\frac{a_{i+1}}{e}, \frac{d_i}{e} * \frac{d_i}{e} * e) \\
&= e * \gcd(\frac{a_{i+1}}{e}, e) && \text{(since } \gcd(\frac{a_{i+1}}{e}, \frac{d_i}{e}) = 1\text{).}
\end{aligned}$$

(b) Before proving part (b) we shall prove (33), (34) and (35) below. Note that $o_h(c_X) = \nu(h) = o_h(a_i) + o_h(b_i)$ for all i. It follows from $k(h) = \nu(h) - o_h(\mathrm{ord}_X(v))$ and $a_1 = \mathrm{ord}_X(v)$, that $o_h(b_1) = k(h)$. We first prove

$$o_h(b_{i+1}) = \begin{cases} 0 & \text{if } o_h(b_i) = 0, \\ 2\, o_h(b_i) & \text{if } 0 < o_h(b_i) \leqslant \nu(h)/2, \\ \nu(h) & \text{if } o_h(b_i) > \nu(h)/2. \end{cases} \tag{33}$$

Suppose first that $o_h(b_i) = 0$. Then $o_h(d_i) = o_h(\gcd(a_i, b_i)) = 0$ and hence $o_h(b_{i+1})$ equals $o_h(b_i d_i) = o_h(b_i) = 0$. This establishes the first part of (33). Next suppose that $0 < o_h(b_i) \leqslant \nu(h)/2$. Then $o_h(a_i) \geqslant o_h(b_i)$, and so $o_h(d_i) = o_h(\gcd(a_i, b_i)) = o_h(b_i)$, which implies that $o_h(b_{i+1}) = o_h(b_i d_i) = 2o_h(b_i)$. Finally, suppose that $o_h(b_i) > \nu(h)/2$. Then $o_h(a_i) < o_h(b_i)$, and so $o_h(d_i) = o_h(\gcd(a_i, b_i)) = o_h(a_i)$, which implies that $o_h(b_{i+1}) = o_h(b_i d_i) = o_h(b_i) + o_h(d_i) = o_h(b_i) + o_h(a_i) = \nu(h)$. Thus (33) is proved.

It is useful to solve the recurrence relation (33). We next prove that

$$o_h(b_i) = \begin{cases} 0 & \text{if } k(h) = 0, \\ 2^{i-1}k(h) & \text{if } k(h) > 0 \text{ and } 1 \leqslant i \leqslant r(h), \\ \nu(h) & \text{if } k(h) > 0 \text{ and } i > r(h). \end{cases} \tag{34}$$

Certainly if $k(h) = 0$ then since $o_h(b_1) = k(h)$ (as we noted above), it follows that $o_h(b_1) = 0$. By (33), we have $o_h(b_i) = 0$ for all i. This establishes the first part of (34). Suppose now that $k(h) > 0$. We next prove (34) for $1 \leqslant i \leqslant r(h)$ using induction on i. The claim in (34) is true when $i = 1$ as $o_h(b_1) = k(h)$. Suppose that $1 \leqslant i < r(h)$ and $o_h(b_i) = 2^{i-1}k(h)$. Then $i + 1 \leqslant r(h)$ and it follows from the definition of $r(h)$ that $2^i k(h) \leqslant \nu(h)$, and hence that $0 < o_h(b_i) \leqslant \nu(h)/2$. Hence by (33) we have $o_h(b_{i+1}) = 2^i k(h)$. Thus (34) holds by induction for $1 \leqslant i \leqslant r(h)$. In particular $o_h(b_{r(h)}) = 2^{r(h)-1}k(h)$. Now by the definition of $r(h)$ we have $2^{r(h)} > \nu(h)/k(h)$, and hence $o_h(b_{r(h)}) > \nu(h)/2$. Hence, by (33), $o_h(b_{r(h)+1}) = \nu(h)$, and by repeated applications of (33), $o_h(b_i) = \nu(h)$ for all $i > r(h)$. Thus (34) is proved.

Equation (34) may be used to compute $o_h(d_i)$. In this paragraph we prove that

$$o_h(d_i) = \begin{cases} 0 & \text{if } k(h) = 0 \text{ or } i > r(h), \\ 2^{i-1}k(h) & \text{if } k(h) > 0 \text{ and } 1 \leqslant i < r(h), \\ \nu(h) - 2^{r(h)-1}k(h) & \text{if } k(h) > 0 \text{ and } i = r(h). \end{cases} \tag{35}$$

Part (a) gives $o_h(d_i) = o_h(\gcd(a_i, b_i)) = \min(o_h(a_i), o_h(b_i))$. Thus if $o_h(b_i)$ equals 0 or $\nu(h)$, then $o_h(d_i) = 0$. This establishes the first part of (35). Consider the second part, and assume that $k(h) > 0$ and $1 \leqslant i < r(h)$. It follows from the previous paragraph that $0 < o_h(b_i) \leqslant \nu(h)/2$. Thus $o_h(d_i) = o_h(b_i) = 2^{i-1}k(h)$ by (34). Finally, suppose that $k(h) > 0$ and $i = r(h)$. By the previous paragraph $o_h(b_{r(h)}) > \nu(h)/2$ and so $o_h(d_{r(h)}) = o_h(a_{r(h)}) = \nu(h) - 2^{r(h)-1}k(h)$ by (33). This proves (35).

The proof of part (b) is now simple. If $k(h) = 0$, then $o_h(b_1) = 0$ and $o_h(a_1) = \nu(h)$ hold. Thus part (i) follows from (34) and (35). On the other hand, if $k(h) > 0$, then (34) and (35) imply that $o_h(b_i) = \nu(h)$ and $o_h(d_i) = 0$ for $i \geqslant r(h) + 1$. Thus part (ii) holds.

(c) We first prove that $s \leqslant r + 1$. Note that $d_i = 1$ is equivalent to $o_h(d_i) = 0$ for all h. Suppose that the number, s, of times that line 5 is executed satisfies $s \geqslant r + 1$. Then it follows from part (b) that $d_{r+1} = 1$, and hence that Is*f*Witness terminates on executing line 6, and $s = r + 1$. Thus $s \leqslant r + 1$. (Note that if $d_i = a_i$ for some $i < r + 1$ then Is*f*Witness terminates at line 7, and $s < r+1$.) Thus $s \leqslant r+1 \leqslant \lfloor \log_2 n \rfloor + 2$ where the last inequality follows as $r = r(h)$ for some h with $k(h) \geqslant 1$ and $\nu(h) \leqslant n$. This proves the second and third sentences of part (c). To prove the last sentence we must show that $r - 1 \leqslant s$ (as $1 \leqslant s$ is clear). This is certainly true if $r \leqslant 2$. Suppose now that $r \geqslant 3$. Fix an irreducible polynomial h such that $r = r(h)$. Then $k(h) > 0$. Showing that $s \nleq r - 1$ is equivalent to showing that Is*f*Witness does not terminate during iteration i when $i < r - 1$. This is equivalent to proving $d_i \neq a_i$ and $d_i \neq 1$ holds for $i < r - 1$ which, in turn, is proved by showing $0 < o_h(d_i) < o_h(a_i)$ for $i < r(h)-1$. The inequalities $0 < o_h(d_i)$ with $i < r(h) - 1$ hold by (35). It follows from (35) and (34) that $o_h(d_i) = 2^{i-1}k(h)$ and $o_h(a_i) = \nu(h) - 2^{i-1}k(h)$. However, $i < r(h) - 1$ implies $i < \lfloor \log_2 \frac{\nu(h)}{k(h)} \rfloor$, which implies $2^i k(h) < \nu(h)$, and hence $o_h(d_i) < o_h(a_i)$. Thus $r - 1 \leqslant s \leqslant r + 1$ and part (c) is proved.

(d) Consider the forward implication. Suppose that Is*f*Witness returns (True, u_s, a_s). Then $d_s = 1$, and by part (a), $\gcd(a_s, c_X/a_s) = 1$ and $a_s \neq 1$. Thus $a_s = \prod_{h|a_s} h^{\nu(h)} \neq 1$, and the Chinese Remainder Theorem gives

$$V(a_s) := \bigoplus_{h|a_s} V(h) = u_s\mathbb{F}_q[X] \subseteq v\mathbb{F}_q[X] \text{ as } V(a_s) \cong \mathbb{F}_q[t]/(a_s) \text{ and } V(h) \cong \mathbb{F}_q[t]/(h^{\nu(h)}).$$

This proves that v is an a_s-witness for X, and X is f-cyclic relative to each irreducible divisor h of a_s. Now consider the reverse implication. Suppose that v is an a-witness for X for some non-constant divisor a of c_X. Then by the definition of an a-witness in Subsection 7.1, for each irreducible divisor h of a, $V(h) \subseteq v\mathbb{F}_q[X]$, and it follows that $k(h) = o_h(b_1) = 0$. Thus by (34), $o_h(b_i) = 0$ and $o_h(a_i) = \nu(h) > 0$ for all i. Hence

$0 = o_h(d_i) < o_h(a_i)$ for all i, and the conditional line 7 of IsfWitness is never executed. It now follows from part (c) that IsfWitness returns True. This proves part (d).

(e) Suppose that IsfWitness returns False, and let h be an irreducible polynomial such that X is f-cyclic relative to h. If $o_h(\mathrm{ord}_X(v)) = \nu(h)$, then $k(h) = o_h(b_1) = 0$, and the argument of the previous paragraph gives that IsfWitness returns True, which is a contradiction. Hence $o_h(\mathrm{ord}_X(v)) < \nu(h)$. Conversely suppose that $o_h(\mathrm{ord}_X(v)) < \nu(h)$, for each irreducible polynomial h such that X is f-cyclic relative to h. Then for each such h, v is not an h-witness for X, and it follows from the previous paragraph that IsfWitness does not return True. Since IsfWitness returns an answer by part (c), it must return False. This proves the first sentence of part (e). The second sentence is an immediate consequence of the first.

(f) The cost of multiplication, division, or finding the greatest common divisor of two polynomials, each of degree at most n, is $\mathrm{O}(n^2)$ field operations. As $c_X(t)$ is an input parameter to IsfWitness, line 3 has cost $\mathrm{O}(n^2)$. Computing $\mathrm{ord}_X(v)$ in line 2 has cost $\mathrm{O}(\mathrm{Mat}(n)\log n)$ by [1, Theorem 6.2.1(b)]. When computing $\mathrm{ord}_X(v)$, one uses "fast spinning" to calculate an $n \times n$ matrix Y with rows $v, vX, \ldots, vX^{n-1}$. We must remember Y in order to compute, for each i, the vector u_i in line 8. If $g_i(t) = \sum_{j=0}^{n-1} g_{ij}t^j$, then $u_i = (g_{i0}, g_{i1}, \ldots, g_{i,n-1})Y$. Thus the cost of lines 8, 9, 10 in the ith iteration of the WHILE loop is $\mathrm{O}(n^2)$. By Theorem 17(c) the WHILE loop is executed at most $\lfloor \log_2 n \rfloor + 2$ times. Thus the total cost of running the WHILE loop is $\mathrm{O}(n^2 \log n)$. Since $\mathrm{Mat}(n)$ is at least $\mathrm{O}(n^2)$, it follows that IsfWitness requires at most $\mathrm{O}(\mathrm{Mat}(n)\log n)$ field operations. □

Remarks. (a) If we use standard algorithms for vector-matrix operations, then an upper bound for the cost of IsfWitness is $\mathrm{O}(n^3)$. For example, at line 2 the cost of finding $\mathrm{ord}_X(v)$ if one uses standard vector-matrix arithmetic is $\mathrm{O}(n^3)$, (see for example, [19, Proposition 4.9]). Similarly, at line 9 we may replace $u := vg(X)$ by $u := ud(X)$. Using the notation of Theorem 17, the sum of the degrees of the polynomials $d_1, \ldots, d_s$ is at most n. Hence the cost of computing $u_1, \ldots, u_s$ is at most $\mathrm{O}(n^3)$. The complexity bound $\mathrm{O}(n^3)$ follows from these observations.

(b) The algorithm IsfWitness may be varied as follows. In essence IsfWitness seeks a divisor a of $\mathrm{ord}_X(v)$ of *maximal degree* satisfying $\gcd(a, c_X(t)/a) = 1$. Although it is straightforward to calculate a from the factorisation of $c_X(t)$ as a product of irreducibles, it is also possible to calculate a using only gcd's and pth roots where p is the characteristic of $\mathbb{F}_q$. We omit the precise details, but the computation of square-free factorizations, see [5, Algorithm 3.4.2], plays an important role.

7.3. **Algorithm IsfCyclic.**

Algorithm. IsfCyclic
Input. a (non-zero) matrix $X \in \mathrm{M}(n, q)$; a positive real number $\varepsilon < 1$
Output. (True, $u, a(t)$), or False

1. $m := \left\lceil \frac{\log(\varepsilon^{-1})}{\log q} \right\rceil$; # m is the maximum number of random vectors tested
2. $c := c_X(t)$; # compute the characteristic polynomial of X
3. $i := 1$; # i counts the number of random vectors chosen
4. WHILE $i \leqslant m$ do
5. $v :=$ a (uniformly) random vector in $\mathbb{F}_q^{1\times n}$; if $v = 0$ then continue; fi;
6. output := IsfWitness(v, X, c);
7. if output $\neq$ False then return output; fi; $i := i + 1$;
8. return False; # probability of failure given that X is f-cyclic is at most ε

Recall that, for an f-cyclic matrix $X \in \mathrm{M}(n, q)$ and a non-constant divisor a of c_X, a vector $v \in \mathbb{F}_q^n$ is an a-witness for X if $v\mathbb{F}_q[X]$ contains $V(a)$.

Theorem 18. IsfCyclic *is a one-sided Monte Carlo algorithm for which, a given matrix* $X \in \mathrm{M}(n, q)$ *and positive real number* $\varepsilon < 1$*, the following hold.*

(a) *If* X *is* f*-cyclic, then* IsfCyclic *returns* (True, u, a) *with probability at least* $1 - \varepsilon$*, where* a *is a non-constant divisor of* c_X*, and* u *is an* a*-witness for* X*.*
(b) *If* X *is uncyclic, then* IsfCyclic *returns* False *with probability* 1*.*

The number of field operations required by IsfCyclic *is* $\mathrm{O}(\frac{\log(\varepsilon^{-1})}{\log q}(\xi_{q,n} + \mathrm{Mat}(n) \log n))$*, where* $\xi_{q,n}$ *is an upper bound for the cost of constructing a uniformly distributed random vector in* $\mathbb{F}_q^n$*.*

Proof. For $m := \left\lceil \frac{\log(\varepsilon^{-1})}{\log q} \right\rceil$, we have $q^m \geqslant \varepsilon^{-1}$. Let X be an f-cyclic matrix relative to at least one irreducible, say h. Suppose that IsfCyclic returns False. Then IsfWitness returns False for m independent uniformly distributed random vectors of V. By the remarks preceding Subsection 7.2, this happens with probability at most $q^{-m \deg(h)} \leqslant q^{-m} \leqslant \varepsilon$, since $\deg(h) \geqslant 1$. If IsfCyclic does not return False, then at least one of the runs of IsfWitness has output (True, u, a), and this is then returned by IsfCyclic at line 7. This proves part (a).

Now suppose that IsfCyclic is has an uncyclic matrix X as input. Then by Theorem 17(e), each run of IsfWitness returns False, and hence IsfCyclic returns False. This proves part (b).

The only situation in which the output of IsfCyclic is incorrect is if the input matrix X is f-cyclic and IsfCyclic returns False. We have shown that this probability of this happening, given that X is f-cyclic, at most than ε. Thus IsfCyclic is a one-sided Monte Carlo algorithm.

Finally, we estimate the cost. Computing $c_X(t)$ in line 2 of IsfCyclic requires at most $\mathrm{O}(\mathrm{Mat}(n) \log n)$ field operations, see [1]. By Theorem 17(f), the total cost of m iterations of the WHILE loop of IsfCyclic is $\mathrm{O}(m\mathrm{Mat}(n) \log n)$ plus the cost of constructing m uniformly distributed random vectors from $\mathbb{F}_q^n$. □

Acknowledgements

The second author acknowledges support of an Australian Research Council Federation Fellowship. Both authors acknowledge the ARC Discovery Project grant DP0879134 which supported the first author's visit to The University of Western Australia.

References

[1] S. Ambrose, *Matrix Groups: Theory, Algorithms and Applications* (PhD Thesis, University of Western Australia, 2006).
[2] G. E. Andrews, *Theory of Partitions* (Reprint of the 1976 original, Cambridge Mathematical Library, Cambridge University Press, Cambridge, 1998).
[3] W. Bosma, J. Cannon and C. Playoust, The Magma algebra system I: The user language. *J. Symbolic Comput.* **24** (1997), 235–265.
[4] R. W. Carter, *Finite Groups of Lie Type* (Wiley Classics Library Edition, Wiley, 1993).
[5] H. Cohen, *A Course in Computational Algebraic Number Theory* (Graduate Texts in Mathematics **138**, Springer-Verlag, 1995).
[6] J. E. Fulman, *Probability in the Classical Groups over Finite Fields: Symmetric Functions, Stochastic Algorithms, and Cycle Indices* (PhD Thesis, Harvard University, 1997).
[7] J. E. Fulman, Cycle indices for the finite classical groups, *J. Group Theory* **2** (1999), 251–289.
[8] J. E. Fulman, P. M. Neumann and C. E. Praeger, *A Generating Function Approach to the Enumeration of Matrices in Classical Groups over Finite Fields* (Mem. Amer. Math. Soc. **176**, no. 830, 2005).
[9] G. Ivanyos and K. Lux, Treating the exceptional cases of the MeatAxe, *Experiment. Math.* **9** (2000), 373–381.
[10] N. Jacobson, *Basic Algebra II* (W. H. Freeman and Co., 1980).
[11] J. P. S. Kung, The cycle structure of a linear transformation over a finite field, *Linear Algebra Appl.* **36** (1981), 141–155.
[12] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.4.10*; 2007, `http://www.gap-system.org`.
[13] S. P. Glasby, The Meat-axe and f-cyclic matrices, *J. Algebra* **300** (2006), 77–90.
[14] S. P. Glasby, A. C. Niemeyer and C. E. Praeger, An f-cyclic recognition algorithm and the Meat-axe, in prep.
[15] S. P. Glasby and C. E. Praeger, Appendix 1: the polynomials $\mathrm{unc}(n, q)$ for $n \leqslant 37$; Appendix 2: Magma computer programs for verifying the conjecture for small n. `http://www.cwu.edu/~glasbys/`.
[16] D. F. Holt and S. Rees, Testing modules for irreducibility, *J. Austral. Math. Soc. Ser.* A **57** (1994), 1–16.
[17] P. M. Neumann and C. E. Praeger, Cyclic matrices over finite fields, *J. London Math. Soc.* **52** (1995), 263–284.
[18] P. M. Neumann and C. E. Praeger, Cyclic matrices and the MEATAXE, in *Groups and computation, III (Columbus, OH, 1999)*, Ohio State Univ. Math. Res. Inst. Publ. vol. **8**, de Gruyter, Berlin, 2001, pp. 291–300.
[19] M. Neunhöffer and C. E. Praeger, Computing Minimal Polynomials of Matrices, *LMS J. Comput. and Math.* **11** (2008), 252–279.
[20] R. A. Parker, The computer calculation of modular characters (the Meat-Axe), in *Computational Group Theory*, Proceedings of the London Mathematical Society Symposium on Computational Group Theory (ed. M. D. Atkinson, Academic Press, London, 1984), 267–274.
[21] R. P. Stanley, Enumerative Combinatorics, vol. II (Cambridge University Press, Cambridge, 1999).

[22] R. Stong, Some asymptotic results on finite vector spaces, *Adv. in Appl. Math.* **9** (1988), 167–199.
[23] G. E. Wall, Counting cyclic and separable matrices over a finite field, *Bull. Austral. Math. Soc.*, **60** (1999), 253–284.

(Glasby) Department of Mathematics, Central Washington University, WA 98926-7424, USA. http://www.cwu.edu/∼glasbys/

(Praeger) School of Mathematics and Statistics, 35 Stirling Highway, University of Western Australia, Crawley 6009, Australia. http://www.maths.uwa.edu.au/∼praeger/